ESD-TR-75-354

MTR-3022, Vol. II

JOBSTREAM SEPARATOR: SUPPORTIVE INFORMATION

JANUARY 1976

Prepared for

DEPUTY FOR COMMAND AND MANAGEMENT SYSTEMS
ELECTRONIC SYSTEMS DIVISION
AIR FORCE SYSTEMS COMMAND
UNITED STATES AIR FORCE
Hanscom Air Force Base, Bedford, Massachusetts

Project No. 522D
Prepared by
THE MITRE CORPORATION
Bedford, Massachusetts
Contract No. AF19628-76-C-0001

REVIEW AND APPROVAL

This technical report has been reviewed and is approved for publication.

Chester G. Clark

CHESTER G. CLARK, Lt Col, USAF
Chief, Engineering Planning Division (MCIO)

Roger R. Schell

ROGER SCHELL, Major, USAF
Techniques Engineering Division

FOR THE COMMANDER

Frank J. Emma

FRANK J. EMMA, Colonel, USAF
Director, Information Systems
Technology Applications Office
Deputy for Command & Management Systems

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: ESD-TR-75-354
2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle): JOBSTREAM SEPARATOR: SUPPORTIVE INFORMATION
5. TYPE OF REPORT & PERIOD COVERED
6. PERFORMING ORG. REPORT NUMBER: MTR-3022, Vol. II

7. AUTHOR(s): J. M. Schacht
8. CONTRACT OR GRANT NUMBER(s): AF19628-76-C-0001

9. PERFORMING ORGANIZATION NAME AND ADDRESS: The MITRE Corporation, Box 208, Bedford, MA 01730
10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: Project No. 522D

11. CONTROLLING OFFICE NAME AND ADDRESS: Deputy for Command and Management Systems, Electronic Systems Division, AFSC, Hanscom Air Force Base, Bedford, MA 01731
12. REPORT DATE: JANUARY 1976
13. NUMBER OF PAGES: 120

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)
15. SECURITY CLASS. (of this report): UNCLASSIFIED
15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

AFWWMCCS
COLOR CHANGE
COMPUTER SECURITY
PERIODS PROCESSING
REQUIREMENTS ANALYSIS
SECURITY KERNEL

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The Jobstream Separator (JSS) has been proposed to automate the manual process used to change security levels at AFWWMCC sites. This volume provides detailed requirements and design analysis information in support of the technical and economic assessment of the JSS presented in Volume I.

ACKNOWLEDGMENT

Project 522D was performed by The MITRE Corporation under the sponsorship of the Electronic Systems Division, Air Force Systems Command, Hanscom Air Force Base, Bedford, Massachusetts.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

SECTION I

INTRODUCTION

This document, Volume II, is composed of a series of appendices that serve as supportive and supplemental data to Volume I of the JSS design. This volume provides additional detail and discussion of the color change process; offers a design of a Honeywell 6000 system Stop/Restart capability; provides engineering data concerning hardware components such as the access control switch apparatus, system display/control panel and machine-readable label processor; offers costing calculations and assumptions underlying JSS cost analysis; and finally, presents a summary of data collection methods and results obtained during the JSS requirements analysis phase of this design effort.

Since each appendix is a separate document, table and figure numbers mentioned refer to only those tables and figures in that appendix.

APPENDIX I

COLOR CHANGE PROCESS

INTRODUCTION

The purpose of this Appendix is to describe the color change process currently in use by WWMCCS sites. The process, as will be described below, is based on several visits to WWMCCS sites (see Appendix VII for data collection method and results). Rather than detail the color-change process used at each WWMCCS site, or enumerate site-specific deviations from the norm, a "generalized composite" color-change process, incorporating the common features from all sites studied, will be treated in detail. This composite process, although not applicable to any one site, will serve as a didactic vehicle for future discussions.

The intent underlying the description and analysis of the color-change process is to:

a. Enumerate all of the steps or subtasks that are performed during the various phases of a color-change.

b. Justify a step's existence -- "Why was this step included in the color-change process?"; "How does this step relate to security regulations?"

c. Analyze the order in which the steps are performed -- "How does one step affect another"? Are steps time dependent? Can concurrent or simultaneous execution of steps be performed? What steps must be performed serially? How long does each step last?

COLOR-CHANGE PROCESS -- OPERATION

Figure I-1 shows a representative daily color-change schedule -- the system is at the unclassified level from 0800-1200 hours, classified from 1230-1600 hours and unclassified, once again, from 1630 onward. The shape of the plot shown is arbitrary and does not represent numerical data collected from any specific WWMCCS site. It serves as a device to emphasize the relative duration of each phase and the relative effect of the color-change process on system utilization and throughput.

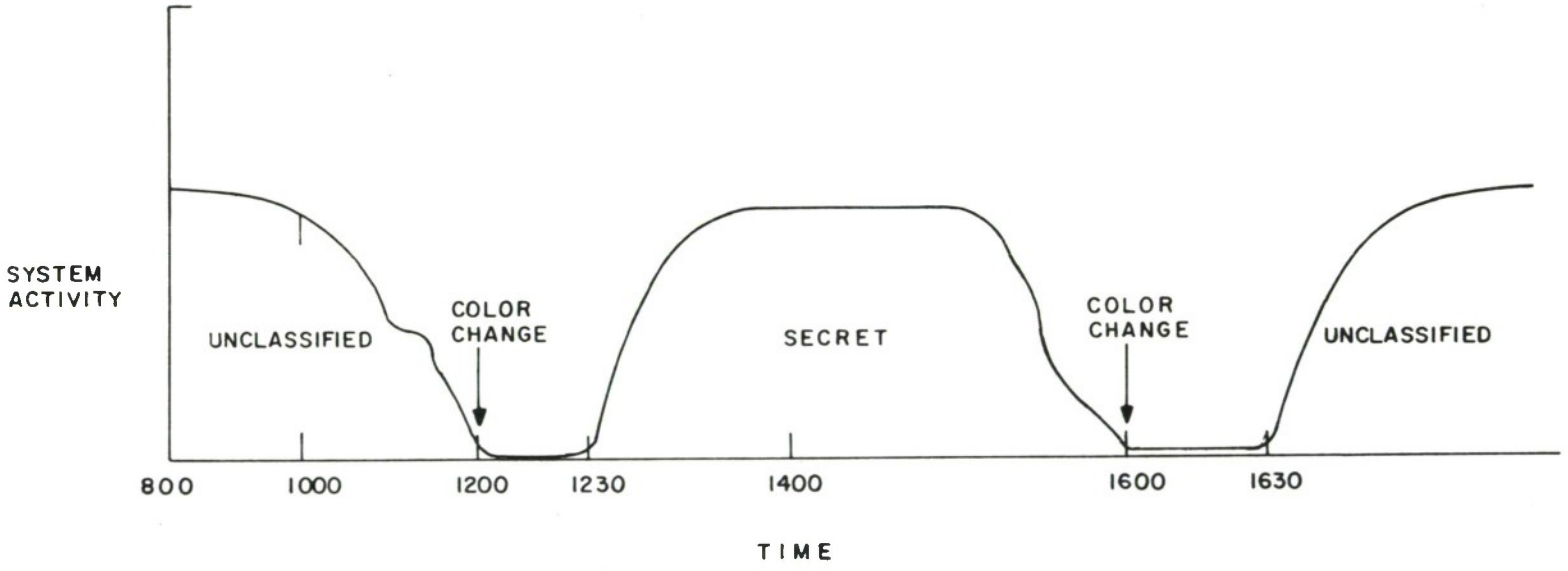

Figure I-1. Typical Color Change Schedule

PHASE-1 -- THE SLOWDOWN PERIOD

The primary function of this phase is to ready the WWMCCS system for the upcoming scheduled color-change. The number of jobs that enter the system is limited at some point prior to the scheduled change. Two advantages are gained thereby; (1) the reduced system workload will enable the system to complete a sizable portion of pending user jobs (if not all), during the time allotted for slowdown; and (2) the reduced workload makes operator scheduling, monitoring, and allocation of system resources an easier task, thereby increasing system efficiency. In order to avoid having to save jobs for processing at a later time when the system comes up again at the same security level, it is beneficial to process as many jobs already in the system as possible, once job entry has been curtailed. The more jobs that are processed during the slowdown phase, the fewer jobs will be required to be saved, maintained, and resubmitted for execution at a later time.

For example, if at the start of the "slowdown" phase there are 20 jobs of relatively uniform requirements (homogeneous mix) awaiting execution in the input jobstream, then it is quite probable that all 20 jobs will be processed and that no jobs will remain enqueued in the system by shutdown time. This ideal situation requires very little operator attention, since the GCOS scheduler assumes full responsibility for the scheduling process.

On the other hand, let us hypothesize a more realistic processing situation that involves a heterogeneous jobstream consisting of (1) small, fast jobs requiring little operator attention and minimal system resources; and (2) huge, multi-step, slow jobs requiring significant operator intervention (i.e., tape mounts) and resource allocation. One would not wish to introduce one of the latter jobs (e.g., a 30-minute job) just before a color change was due (e.g., 10 minutes prior to the change). If processing is allowed to occur on an "as usual" basis without special attention from an operator monitoring the input queue, the system scheduler will eventually permit one of the large, slow jobs (should one exist) to begin execution. The possible results of such a job entering the execute mode are:

a. System throughput drops precipitously, due to the monopolization of system resources by the big job.

b. The upcoming color change may have to be temporarily delayed until the big job finishes execution. The delay experienced will either cause the duration of the next security level to be shortened in order to maintain the daily color change schedule (i.e., a 1200 hour change would be delayed until 1300 hours, but the new color would be 1 hour shorter in duration), or will cause a fixed delay in the start of all successive periods (i.e., the 1200 and 1800 periods will each start 1 hour later).

c. The big job may be summarily aborted in mid-execution if it is still executing at the time for the scheduled color change. In this case, the big job's spoolout would be lost, the job would have to be resubmitted for execution at a later time, the user would lose significant turnaround time, and all processing and allocation of system resources up to the point of termination would have been for naught.

d. The big job may finish execution prior to the scheduled color change time, but due to the lack of time remaining in the slowdown period, "smaller" jobs still awaiting execution may not be executed, and would be forced to suffer additional delays until processing resumes at the same security level later in the day. Thus, the big job can significantly affect the turnaround times of other jobs. It also creates additional work for the operator, who must now save the input job queue (e.g., on tape), and later resubmit (via tape and/or cards) the unprocessed jobs in order to recreate the "leftover" job queue.

It is essential that the operator monitor the input job queue during the slowdown phase, in order to facilitate the emptying of that queue and reduce the likelihood of impact on system performance and user turnaround time. A prudent, watchful operator can significantly effect Phase-1 processing activity and duration.

The physical removal of media and setting/resetting of devices are a few of the major functions of the operator. The operator also plays a key role during Phase-1 job monitoring and selection, as certain jobs in the input job queue are processed rather than others in order not to impact "slowdown" activity. The operator must base his job selection on user-supplied estimates concerning a job's characteristics (i.e., CPU time, I/O activity, and memory usage). These estimates, which are routinely attached by the user to run-time submission sheets that accompany the input job, can knowingly, or unknowingly, be in error. There is no way the operator can insure that a job selected for execution on the basis of time and

resource estimates supplied by the user does match the estimates until the job is actually executed.

The likelihood that jobs that were submitted during normal processing will still remain enqueued by the end of the slowdown phase is quite high, depending upon factors such as operator awareness and experience, system workload and system configuration. Using the example shown in Figure I-1, we see that unclassified jobs entered at 1000 hours, but still not processed by 1200 hours, must be resubmitted during a later unclassified processing session (1600 hours). These "old" jobs, which have aged from 6 to 8 hours, must now compete with "newer" unclassified jobs entered during the 1600 session. Thus, the turnaround times associated with such jobs serve as a clear indicator of the inadequacies of period processing, and color change procedures in general. Not only are jobs delayed, but the overhead in physically maintaining media associated with these jobs adds to operator workload and therefore increases the possibility of a breach in security.

Spoolout data generated during normal processing presents a delay during slowdown. The SYSOUT data file, which contains all spoolout data, must either be emptied prior to system shutdown (as is done at SAC) or must be saved and processed during a later session (as is done at MAC). The decreased workload on the processor during the latter stages of slowdown permits increased I/O activity as CPU contention and system workload decrease and availability of system resources increases. This factor contributes to the emptying of the SYSOUT file but, often, this increased availability is not enough. Depending upon the system configuration (i.e., number of printers, punches, etc.), the status of each unit (operable vs. inoperable), and the quantity of spoolout activity, the time required to print the entire file may be excessive (at SAC, for example, draining of output queues may take as much as 1 to 2 hours in some instances). User jobs that have long since finished execution may have generated large quantities of printout that is still awaiting printing. Thus, a decrease in the number of jobs in the system does not immediately lead to a decrease in spoolout activity. In fact, a time plot of job activity versus spoolout activity would show spoolout "lagging" job activity as time for the color change approached.

The previous discussions focused on the "batch oriented" aspects of color changing. Equal attention must, however, be given to the interactive user whose role is also significantly affected by the period processing. Thus, Phase-1 activity must account for the interactive user population at remote terminals. Although the user is generally aware of the daily processing schedule, the system

operator broadcasts a warning message to inform all users of an impending change, and issues an NCALL system command, which prohibits new users from logging on the system. The reminder or warning message usually occurs 20 to 30 minutes prior to the scheduled changeover (estimate obtained from MAC, SAC, and DSDC). In certain circumstances, a courtesy phone call to special users will precede any action taken by the system operator. Once having received the message, the user is given approximately 5 minutes to tie up all loose ends, save his files and logoff. At the end of this 5-minute period, the operator will summarily abort all interactive users who have not as yet logged-off by issuing a TCALL system command. Thus, all interactive processing is curtailed approximately 15 to 20 minutes prior to the scheduled color change time.

Phase-1 Checklist

Table I-I and Figure I-2 summarize the steps performed during Phase-1. Table I-I offers a brief description of each step while Figure I-2 depicts chronologically the order and duration of each step. The encircled numbers on Figure I-2 refer to the step numbers in Table I-I.

The events listed in Table I-I and the step durations and general shape of the system throughput plot are approximations based on statistics collected at Military Airlift Command (MAC) and Strategic Air Command (SAC) WWMCCS sites (see Appendix VII for description of data collection method and summary of trip reports).

PHASE-2 -- THE CHANGEOVER PERIOD

To insure the complete separation of security levels, so that users at the new security level can in no way access user or system information that was processed at the previous level, it is necessary to erase all remnants of programs and data left in core, to dismount all storage media, and to clear all buffers and memories in various peripheral equipment. Typical system devices affected are: the DATANET 355 communication control processors (DN 355), System Controller (SC) units, Input/Output Multiplexor (IOM), Microprogrammed Peripheral Controllers (MPC), etc. Once these steps have been taken, the first half of Phase-2 activity is complete and a sterile sanitized environment exists. The second half of Phase-2 activity requires the team of operators to perform a series of manual operations in order to create the new security level.

Table I-I

Summary of Phase-1 Activity

Step	Event
0	Baseline activity level -- average processing workload and system throughput.
1	Color change preparation begins. Operator monitors batch job activity -- no more new batch jobs are accepted. At some pre-specified time (site-specific) prior to scheduled change-over, operator begins to monitor the input job queue to select jobs that will not impact period processing schedule. Screening process (if employed) usually begins about 45 minutes to 1 hour prior to color change.
2	Warning message is sent to all interactive users noting impending color change at 1200 hours. User has 5-minutes to save files and logoff. The warning message is usually sent approximately 20 to 30 minutes prior to the scheduled color change.
3	Operator issues NCALL system command to prevent additional users from logging onto system.
4	All interactive users must logoff: TCALL issued to sever all remote users. In certain instances, another warning message or telephone call will be made to find out what is causing delay in user logoff. TSS (time-sharing) is terminated and the communication lines and/or modems are detached.
5	Spoolout finishes -- if not, save-tape is made for later processing.
6	System and user file backup procedures (tape-to-tape, disk-to-tape) are performed. These activities are not usually considered to be part of the color change process, but are included here to indicate the type of site-specific, system-related, activities that are often performed during Phase-1 processing.
7	Phase-1 activity is complete -- Phase-2 may begin. The status of the system at this point is as follows: - All interactive users have logged off. All batch jobs have either been run, or will be resubmitted for execution during a later period. - Spoolout activity has ceased -- spoolout has either finished completely, is being processed off-line, or will be processed by the Honeywell system during a later period. - System and user file back-up procedures have completed. - Only GCOS system modules are active -- the core allocator and peripheral device allocator mechanisms.

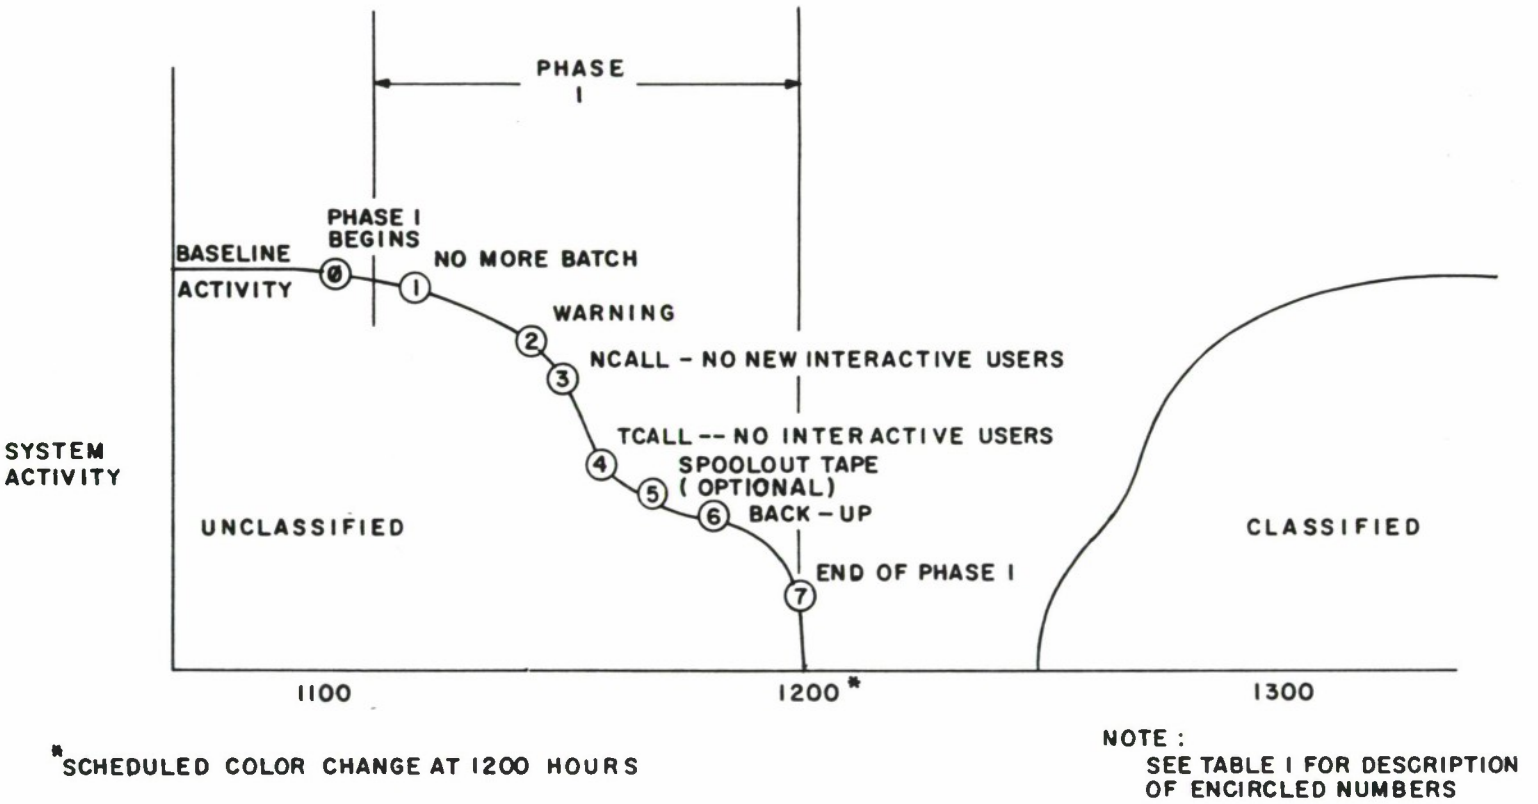

Figure I-2. Chronicle of Basic Phase-1 Activities

Operator intervention is considerable during Phase-2 activity, since most of the actions require manual procedures (i.e., mount-dismount, attach-detach, set-reset, etc.). Although as many as 5 or 6 operators are performing specialized tasks (i.e., one operator performs tape mounts, another handles console messages), the overall activity during Phase-2 requires approximately 30 minutes for the completion of assigned tasks. This time estimate varies depending upon three factors: operator performance (which is usually quite good), system configuration, and in some instances, the relative "direction" of the color change (i.e., SECRET→TOP SECRET or SECRET→UNCLASSIFIED, etc.).

Even though there are differences among large and small WWMCCS sites, in terms of system configuration, workload, and mission (which determines the "colors" used and therefore the number of color changes to be made), the 30-minute Phase-2 estimate is, for the most part, a WWMCCS system constant. The site-to-site deviation associated with this Phase-2 time is approximately 5 minutes. In other words, no matter how large or small the WWMCCS site is, Phase-2 duration will be between 25-30 minutes.

The security standards imposed on the operators in the tasks they perform impose redundancy to assure that no security errors occur. Operators are burdened with a time-consuming "check-and-double-check" policy. For instance, operator A sets a switch in a device and when he is finished, operator B checks the same device to see if operator A has done it correctly. Other time-consuming security related policies range from the minor delay incurred when operators physically change their dress code (operator's jacket color signifies level of current operation), to the segregation (isolation) of storage media of different classifications.

Phase-2 Checklist

The operations performed during Phase-2 of the color change are enumerated in this section. Since the steps performed vary depending upon the relative direction of the change (i.e., Classified to Unclassified or vice-versa), two separate checklists are included.

Table I-II lists the 16-step checklist used in changing from Unclassified to Classified levels. The steps performed are chronologically depicted in the 1200 hour color change in Figure I-3.

Similarly, Table I-III lists the 20-step checklist used in changing from Classified to Unclassified levels. The steps performed are chronologically depicted in the 1600 color change in Figure I-3.

Table I-II

Phase-2 Unclassified to Classified Checklist

Step	Event
1	The "BOOT" system command is entered at the operator's console. The system responds with a question relsting to the source of the boot program. The operator typically replies with "CARD" to indicate the bootload source. The BOOT command places an end-of-file (EOF) mark on the system accounting tape.
2	The following items are removed and stored: a. All disks and tapes b. All cards (keypunch, reader, punch) c. All listings
3	All CRT's in machine room are cleared.
4	To preclude mixing of different level material, the following steps are done: a. Tape carts are moved to unclassified section of room as are listings, waste, and unfinished products. b. A double check procedure is begun before bringing in the classified data. c. Signs are put up to indicate that classified processing has begun.
5	Operators change jackets to signify new processing level.
6	The following items are obtained from library: a. Classified disk packs b. Pack save tapes c. Scratch tapes d. Ribbons and stamps e. Number for next classified waste-bag
7	Classified "PERM" packs and data base packs (and tapes) are mounted.
8	Classified accounting tape is mounted.
9	Console and printer ribbons are changed.
10	Punches are loaded with new colored cards.
11	CRT's are locked out, if not already done.
12	Reboot using INIT and BOOTLOAD switches on console.
13	Answer "Restart" and "System Scheduler Restart" questions at console.
14	Boot the micro-programmed controllers (MPC) -- disk and tape control units.
15	Check classified storage cabinet for any jobs previously submitted. These jobs are run first.
16	The DATANET communication control computers are BOOTLOADED.

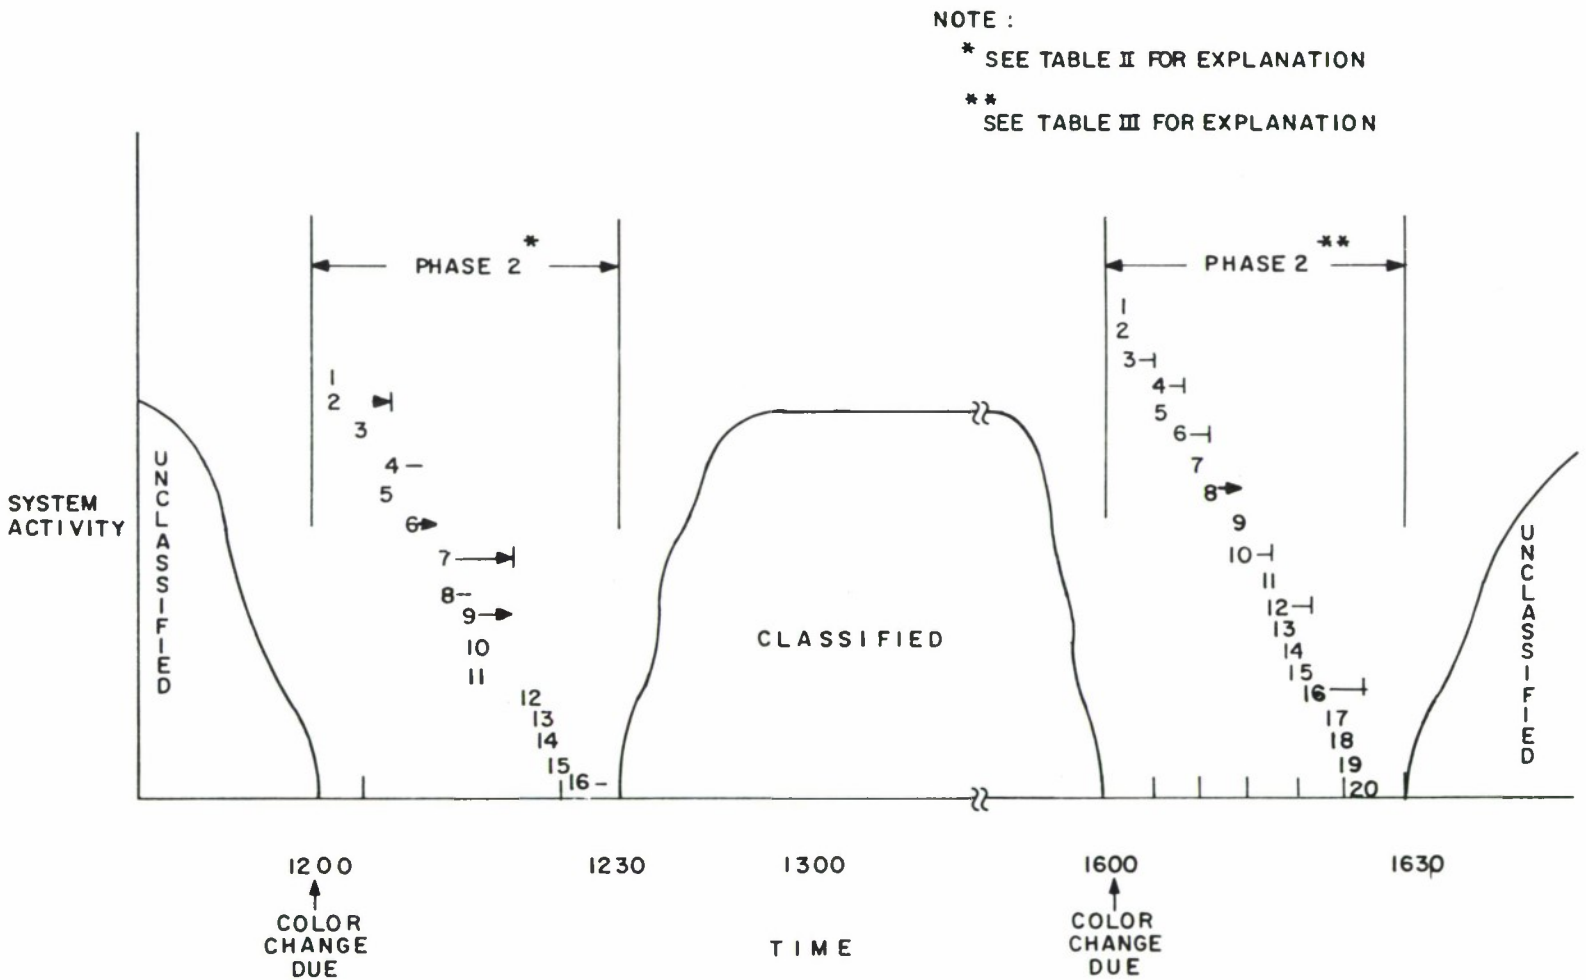

Figure I-3. Chronicle of Basic Phase-2 Activities

Table I-III

Phase-2 Classified to Unclassified Checklist

Step	Event
1	Remove cards from punch.
2	Clear face of all CRT's.
3	Run program to clear buffers and registers in DATANET.
4	After conclusion of DATANET clear program, run a dump program that prints out contents of DATANET memory.
5	Mount a new scratch tape.
6	A 1-card clear core program is loaded in the system card reader, and the "INIT/BOOT" botton is depressed.
7	Stop all disk drives, initialize the MPC's (which serve as their controllers), and dismount the packs.
8	A 6-card dump program is placed in the card reader and run when the INIT/BOOT button is depressed. The dump program prints out the contents of core containing non-zeroes. The security office is informed if a problem arises with the core-clear program (i.e., specific locations are not zeroed out properly).
9	The console listings, 355 listings and dumps, are collected, dated, marked with the highest classification and stored.
10	The ribbons from printer and console are removed and stored. Cards are removed from the reader, punch and keypunch.
11	Dumps, console sheets and checklists, disk packs, classification stamps and unused cards are stored in the classified cabinets, SECRET waste is disposed. Remove jackets.

SECRET level has been annihilated -- UNCLASSIFIED level can now be created.

Step	Event
12	Mount unclassified ribbons.
13	Get boot tape and latest save tapes.
14	Reconfigure the system.
15	Turn bus-back selection switches of modems to enable connections to remote sites.
16	Mount unclassified perm packs.
17	Reboot system using INIT and BOOTLOAD switches.
18	Answer "Restart" and "System Scheduler Restart" question.
19	Bootload the micro-programmed controllers.
20	BOOTLOAD DATANET communication controllers.

PHASE-3 -- THE POST COLOR CHANGE PERIOD

During Phase-3, system operation restarts at the new security level. Figure I-4 chronicles basic Phase-3 activity. Although only 4 or 5 minutes are required to perform the Phase-3 steps, Figure I-4 shows Phase-3 duration to be about 15 minutes -- the time needed for the system to become fully operational (all user options made available) and "base-line activity" achieved.

Figure I-4 indicates a slow increase in system throughput as would be expected, since the operating system has just been loaded and the jobstream is non-existent. Once the various GCOS system modules that control special user options (e.g., time-sharing) are loaded, new and "old" jobs can be submitted for processing. Any "old" jobs submitted for execution during a previous period of the same classification can now be resubmitted. Usually these jobs receive higher priority in the input queue than "newer" jobs submitted either during Phase-2 or at the start of Phase-3.

CRITIQUE OF CURRENT COLOR-CHANGE PROCESS

This section presents a discussion of the effects of the color change process, as seen through the eyes of those who are most affected by it; namely, the operators, the users (batch and interactive), the scheduling staff, and most important, the outside agencies that depend upon the WWMCCS site to perform its basic command and control mission during times of crisis.

The User's View

Regardless of classification or mode of processing, the WWMCCS user suffers directly and indirectly from the effects of period processing at various security levels. Some effects are:

a. Early logoff of interactive users.

   The user must contend with a "shortened" work period that is a direct result of security-related color change policy. Once the NCALL command is issued, new interactive users cannot logon to the system. Thus, system availability to the user is shortened by about 30 minutes per color period.

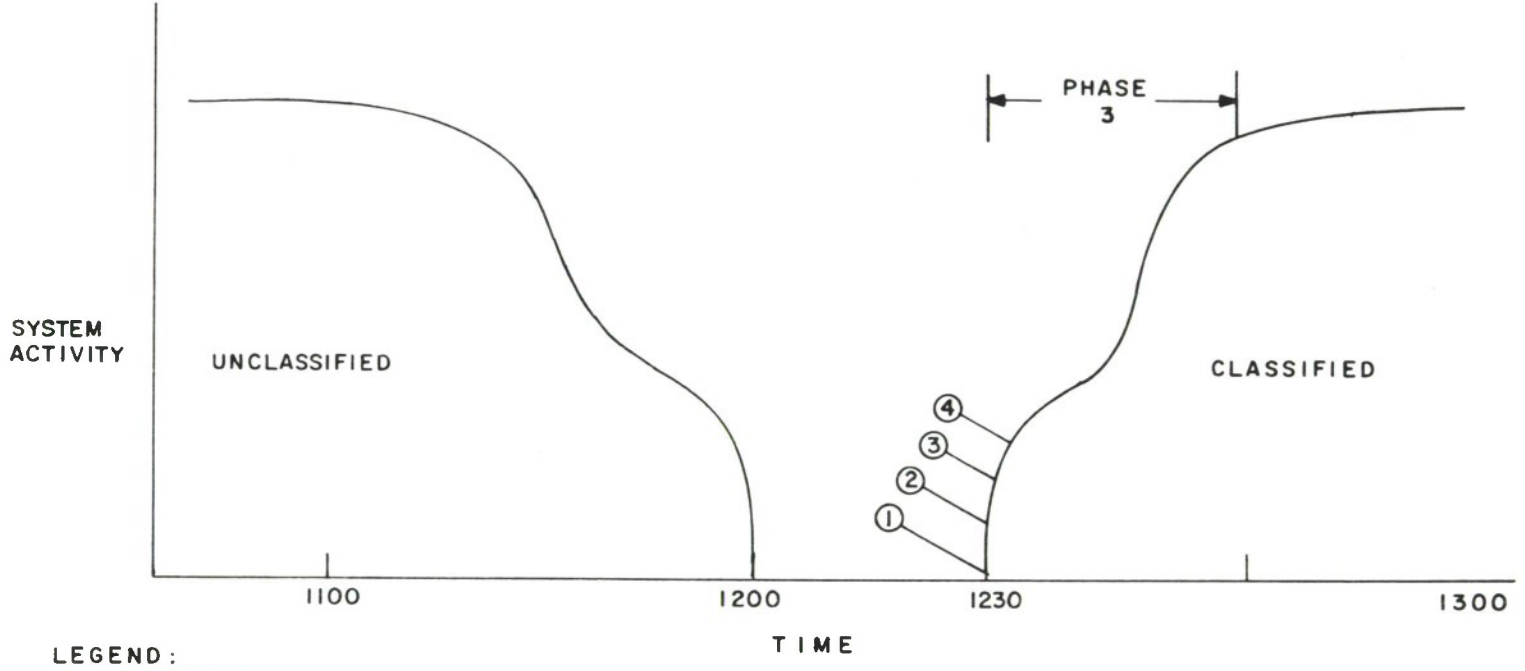

Figure I-4. Chronicle of Basic Phase-3 Activities

b. Turnaround time.

   Depending upon the color change schedule and the type, size, and time of submission, the turnaround time for a MAC or SAC job could, as a worst case, range from 5 to 7 hours.

c. Processor unavailability during Phase-2.

   Since no system activity is permitted during Phase-2, the user community is completely denied about 30 minutes of processing time per period.

d. Excessive wait between desired security levels.

   The excessive waiting experienced by users may have a definite negative effect on user productivity. Once the user's desired color becomes active, the user must perform on a "get-as-much-done-as-possible" basis, since he is faced with the prospects of having to wait several hours before processing will resume at the current level. The realization that a finite duration processing period will be followed by a rather lengthy hiatus impacts the quality and quantity of work accomplished by the user in the allotted time frame. The psychological burden on the user, although immeasurable, contributes to a decrease in productivity with concomitant increase in the probability of stress-induced errors.

ADP Scheduler's View

Period processing not only affects users, who must work within the time bounds of prescheduled security levels at fixed periods, but also complicates the function of the ADP staff responsible for the preparation of a reasonable schedule in the first place. The inability to dynamically schedule processor usage affects productivity and system utilization. Most sites, at present, schedule system usage at least one week in advance. The schedule is based upon basic mission requirements, daily or weekly preventive maintenance requirements, software development work, past workload statistics for the same period, and projected workload statistics for the upcoming week. Once these factors have been weighted according to their priority, a schedule is created that will be stringently adhered to during the upcoming week. Thus, the schedule created is based on mission and user requirements that will be a week old by the time the schedule is in operation.

If the schedule does not reflect the needs of the user community, but favors the operations staff in that only a minimum number of color-changes need be performed per day, the obvious result is general user dissatisfaction as some mission tasks are not given enough system time. On the other hand, if the schedule tries to serve "too many masters", by calling for short periods and numerous changes in security level, then the ADP operations staff would suffer the consequences. Manual intervention would increase drastically as operators would be required to sprint about the machine room performing endless color-change actions. Short periods not only physically exhaust the operators, but drastically affect system utilization and throughput. The system state during such periods would be in one of three categories: "going-down", "down", or "about to come-up". Users would barely have enough time to restart their activities when preparation for the impending color change would necessitate curtailment of their activity.

The difficulty in scheduling arises in clustering the various users into processing periods that satisfy all parties. If two Secret user applications must be run at different, but specific times of day, then there is no alternative but to schedule at least two Secret periods.

ADP Operator's View

The ADP staff, who bear the brunt of the color-change procedure in terms of time and energy spent, must perform the ritual as efficiently and accurately as possible, regardless of the site's mission, job workload (backlog), system configuration, security level, or operational mode (crisis versus daily processing).

Operators at MAC and SAC perform about 50 tape and 50 disk mounts per day in satisfying user and GCOS I/O requests that are independent of color-change requirements. The color-change process increases these figures substantially. All permanent packs (i.e., those that contain parts of the GCOS operating system) in addition to the data base (packs and tapes) must be dismounted and remounted for each change in security level -- a total of 45 mount and 45 dismount operations per day at MAC, for example. These figures were derived by assuming that 10 disk and 5 tape drives, or 15 drives in all, must be dismounted and mounted -- a total of 30 operations per period. If 3 periods are active per day, then 90 operations are performed that directly relate to color change activity.

In light of the site-specific security requirements in effect, operators must not only be cognizant of their own actions, but are

also responsible for the actions of other operators. This policy results in the use of "check and double check" procedures instituted to safeguard all manual intervention (i.e., mounting/dismounting storage media, clearing and resetting of memories, buffers, and registers in the SC's, MPC's, and IOM's, etc.).

WWMCCS Command and Control Function

The delays incurred by period processing via the color-change process seriously hamper the execution of command and control functions at WWMCCS sites. Any delays encountered at an individual WWMCCS site will be felt by those agencies, or other sites, that depend upon that site for processing related to command and control functions.

A less obvious, but extremely important, drawback to current color-changing procedures, in light of the WWMCCS command and control mission, arises whenever a world-wide military crisis or emergency situation occurs. In past cases, such exigencies have necessitated the immediate processing of classified programs and data in order to fulfill mission tasks at various WWMCCS sites.

Since a "slowdown or run-out period" cannot be performed, jobs currently in execution, whatever their size, priority, or status, are summarily aborted in favor of crisis-handling programs.

The color-change introduces sufficient delay time into the WWMCCS command and control mission to reduce the usefulness, effectiveness and applicability of certain programs and data designed to handle such crises. During the half hour needed for security level change, the nature of the crisis could change so drastically that the classified programs intended for execution at the new level would no longer be applicable to the situation at hand. At present, sites playing significant roles in the WWMCCS command and control mission must forego unclassified processing during crisis situations, because these sites lack physical security precautions necessary to insure that users at remote terminals accessing unclassified data bases in a host machine are non-malicious and are of the appropriate security level. Tapping of input/output lines, by direct access or by means of electromagnetic interception, prevents the use of sites that maintain large on-line data bases that serve a multitude of remote users during crisis mode.

APPENDIX II

PROPOSED STOP/RESTART FACILITY FOR HONEYWELL 6000 SERIES

INTRODUCTION

This appendix discusses three GCOS-related problems associated with JSS operation:

1. quiescing the Honeywell 6000 computer system that is processing data at a particular classification,
2. clearing the computer system, and
3. restarting the computer so that data at a different classification level can be processed.

Quiescing consists of stopping all processors (central processing units and I/O processors) at a point from which they can be restarted at a later time, and saving the contents of the H6000 system's memory and registers. Clearing the computer system involves initializing all memory and registers, so that data of one classification level is not present within the system when programs and data having a different classification level are entered. Restarting the computer consists of restoring the contents of H6000 memory and registers and allowing processors to continue from the point at which they had been quiesced.

APPROACHES

Two approaches to quiescing/restarting are presented here (see Figure II-1). One approach is directed towards quiescing/restarting at the interface between the H6000 hardware and the GCOS Hard Core Monitor (HCM). The HCM is that collection of GCOS software modules that comprise the nucleus of the operating system. The second approach consists of quiescing/restarting at the interface between the GCOS HCM and GCOS privileged slaves. The two approaches are discussed separately and compared in later sections.

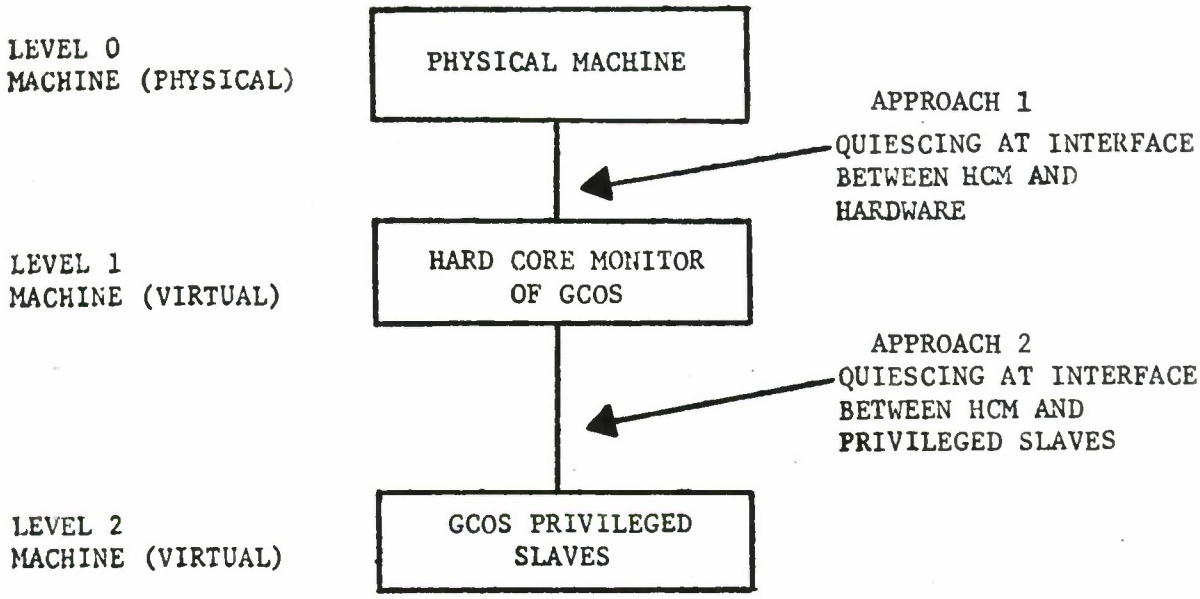

Figure II-1. Two Approaches to Quiescing/Restarting

The ensuing discussion will focus on the quiesce/restart process, whose function is based on fulfilling five objectives:

A. Activate quiescing software in the H6000 and put central processing units into a disabled state.

B. Stop all I/O operations between the Input/Output Multiplexor (IOM) and H6000 memory. Quiesce the DATANET 355 communication control computer.

C. Perform any operations necessary for later device repositioning, save the contents of H6000 memory, and signal the minicomputer that quiescing has completed.

D. Disconnect data paths to devices for the current level, clear H6000 memory and registers, and connect devices for the new level.

E. Restart the system so that processing at a new classification level can begin.

APPROACH I: QUIESCING AT THE INTERFACE BETWEEN GCOS HCM AND H6000 HARDWARE

This approach can be implemented by adding a routine to the GCOS Hard Core Monitor (see Figure II-2). The instructions of this routine are executed when a processor receives the quiesce request through a fault or interrupt. The following sections indicate how each of the five objectives may be met. Table II-I summarizes Objectives A-E for Approach I and provides timing estimates for each objective.

Objective A

This objective requires that the H6000 quiescing software be activated and that the processors (other than processor 0) be put into a delay-until-interrupt state. Processor 0 will be used to execute the instructions for stopping I/O (objective B). Before discussing methods for fulfilling objective A, it is appropriate to present some background information on the H6000 fault and interrupt mechanisms. There is one fault vector per processor. There is one entry in a fault vector for each fault type. Four of the fault types are particularly relevant to quiescing - connect, execute, startup, and shutdown. The connect fault is a means by which one processor can signal (i.e., cause a fault in) another processor. The startup, shutdown, and execute faults are generated by manually

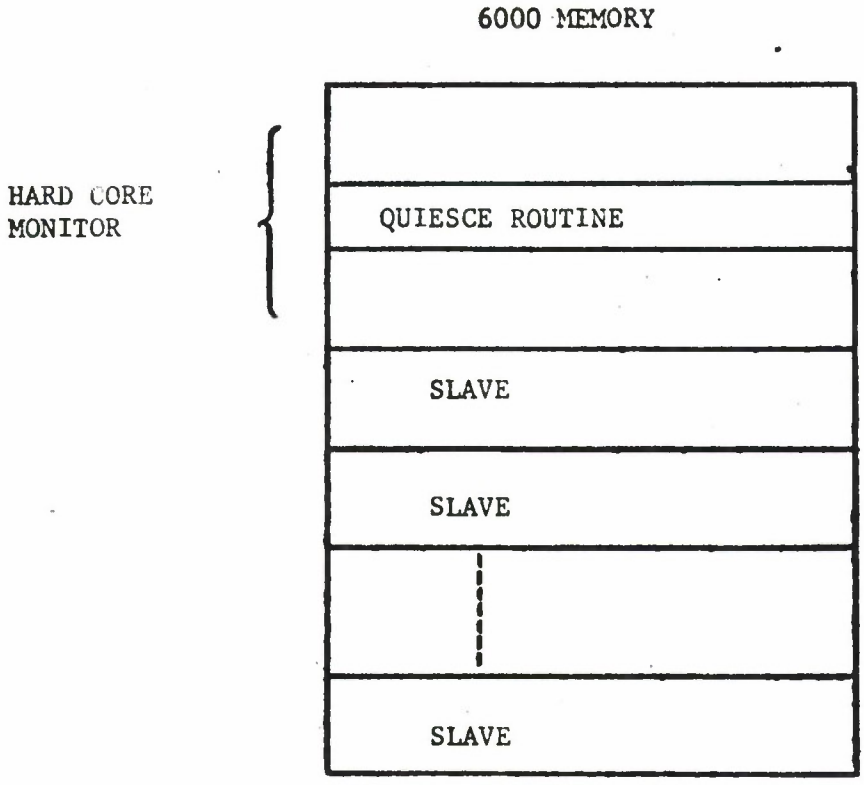

Figure II-2. Quiescing Between HCM and Hardware

Table II-I

Time Required to Quiesce at Interface Between
H6000 Hardware and GCOS Hard Core Monitor

OBJECTIVE	TIME REQUIRED	COMMENTS
A	milliseconds	Insignificant
B	7 minutes	The time required here is bounded by the longest possible I/O operation. The time required by an I/O operation is the lasped time between the CONNECT to an IOM and the interrupt associated with the end of the I/O operation. A tape forward file space command which must travel the length of the tape with the read/write head down could take seven minutes. This is considered to be the longest realistic I/O operation.
C	2 minutes	The time required to fulfill this objective is determined by the maximum amount of time required to rewind a tape drive that is being shared between classification levels. The time required to dump 6000 memory to disk is not significant. With a DSS-190, one cylinder (55,000 words) of data can be transmitted in 325.4 millisecs. This means that one quadrant of memory (256K words) can be transmitted in about 1.6 secs.
D	30 seconds	If the clearing of memory and registers can be done by the minicomputer and 6000 purge program (without operator intervention), this time should be insignificant. An upper bound should be about 30 seconds. This allows time for the operator to place the 6000 purge program deck into the card reader.

Table II-I (Concluded)

E	4-6 minutes	The time required here is dependent upon the operators mounting disks and tapes, putting cards in card readers, changing printer ribbons, and repositioning paper in printers. Operators should require 30 seconds to mount a disk pack. Once a DSS-180 pack has been mounted, it requires about 1.5 minutes to power-up and become ready (DSS-190 disks require only about 30 seconds). Powering-up can be overlapped with other operations. The time for objective E will depend upon the number of operators, number of shared devices, and the operational procedures. NOTE: If the restart software must reposition tapes before the 6000 system is allowed to begin executing, then the time for objective E could be as much as 8-9 minutes.

pushing a button on the processor maintenance panel. If a processor receives a request for a startup or execute fault, it may respond to that request (i.e., the processor begins executing the instructions contained within the startup or execute entry of this processor's fault vector) before it finishes with the instruction that it is currently executing. A processor will finish the instruction that it is currently executing before responding to a request for a shutdown fault.

There is one interrupt vector for each Input/Output Multiplexor (IOM). There is one entry for each of four interrupt types - special, initiate, terminate and marker - within an interrupt vector. When GCOS software is running on the H6000 hardware, switches on the processor and system controller maintenance panels are manually set so that only processor 0 will respond to interrupt requests.

Two basic strategies could be used to accomplish objective A. In the first strategy processor 0 responds to the initial signal to quiesce and it then signals the other processors (see Figure II-3).

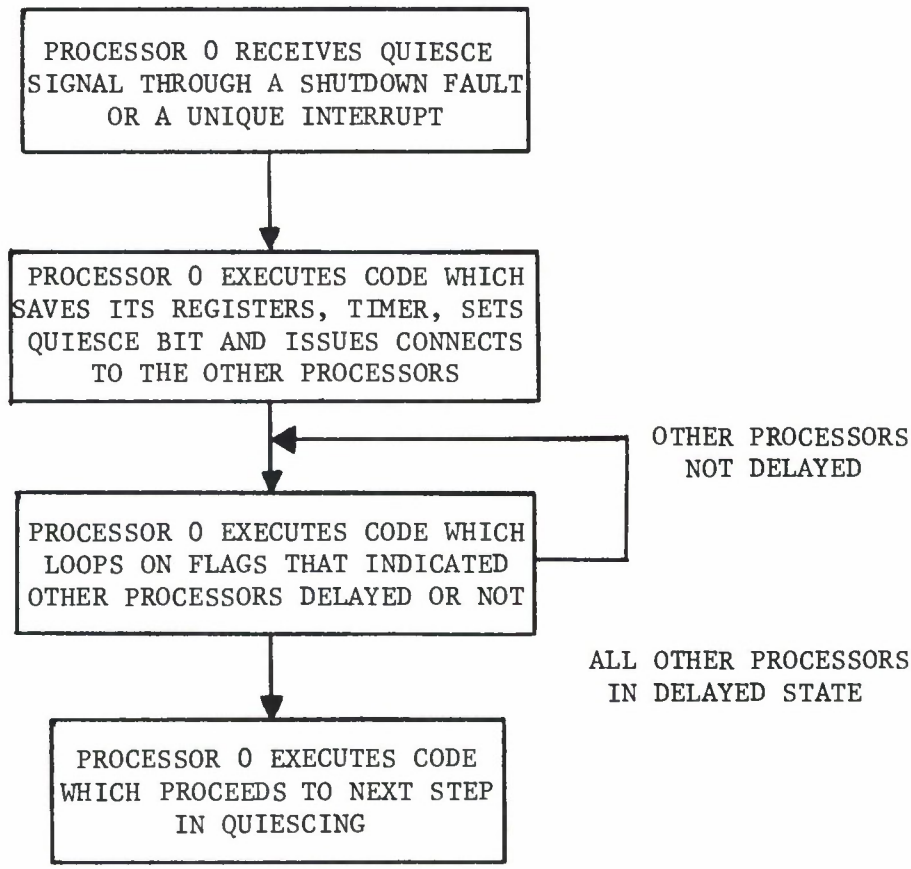

Figure II-3. Quiesce Routine Executed by Processor 0 (Strategy 1)

The initial signal could be sent to processor 0 by means of a fault or interrupt. The shutdown fault would be an appropriate initial signal since processor 0 will respond to this fault only when this processor has completely finished executing the current instruction. This means that at some later time processor 0 could return to executing instructions at this point and the resulting computation would be equivalent to a computation in which processor 0 executed these instructions without responding to the shutdown fault. The shutdown fault entry in processor 0's fault vector would contain the address of a software module which recognizes the occurrence of a shutdown fault as being the initial signal to start quiescing. The initial signal could also be an interrupt. For example, the convention could be adopted that the occurrence of a special type interrupt on payload channel* number 5 of IOM number 0 is to be interpreted as the initial signal to quiesce. The entry for special type interrupts in the interrupt vector for IOM number 0 would contain the address of a software routine which recognized a special interrupt on payload channel number 5 as being the quiesce signal. If the minicomputer is interfaced to the H6000 by means of an IOM payload channel, then the interrupt representing the quiesce signal could be generated from that IOM and payload channel to which the minicomputer is connected. After responding to the initial quiesce signal (shutdown fault or interrupt) processor 0 executes code to store and save its registers and then issues a quiesce signal to the other processors. One method for signalling other processors would be for processor 0 to execute a CIOC instruction causing a connect fault to occur in each of the other processors. Figure II-4 indicates the logic of the existing connect fault routine within GCOS. The connect fault is currently used to allow one processor to remove another processor from a delay-until-interrupt state. Figure II-5 indicates how the logic of the connect fault routine could be modified so that the occurrence of a connect fault is also interpreted as being a quiesce signal.

Another strategy for fulfilling objective A would be for all processors to respond to the initial quiesce signal. This could be done by generating a shutdown fault in every processor. Figure II-6 indicates the logic that could be incorporated into the shutdown fault routine so that the occurrence of this fault is interpreted by all processors as being a quiesce signal. At the present time the

---

*A payload channel is a data path between an IOM and a peripheral controller.

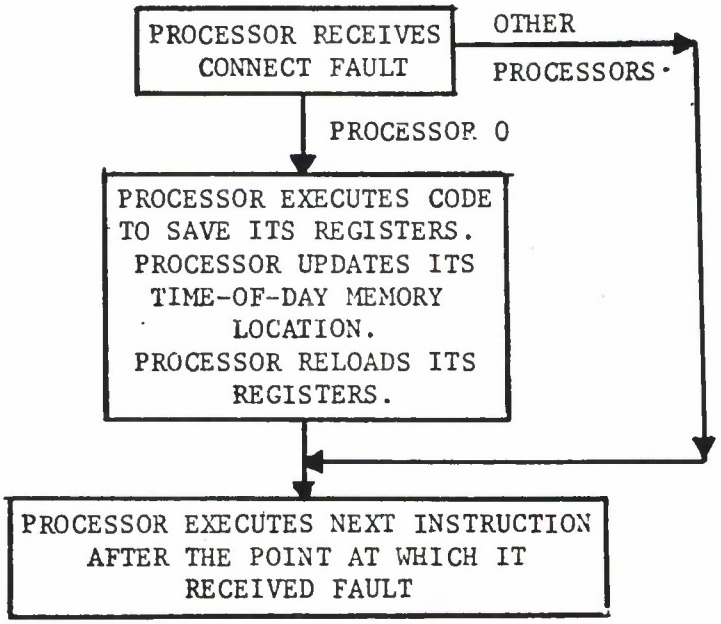

Figure II-4. Existing Connect Fault Logic

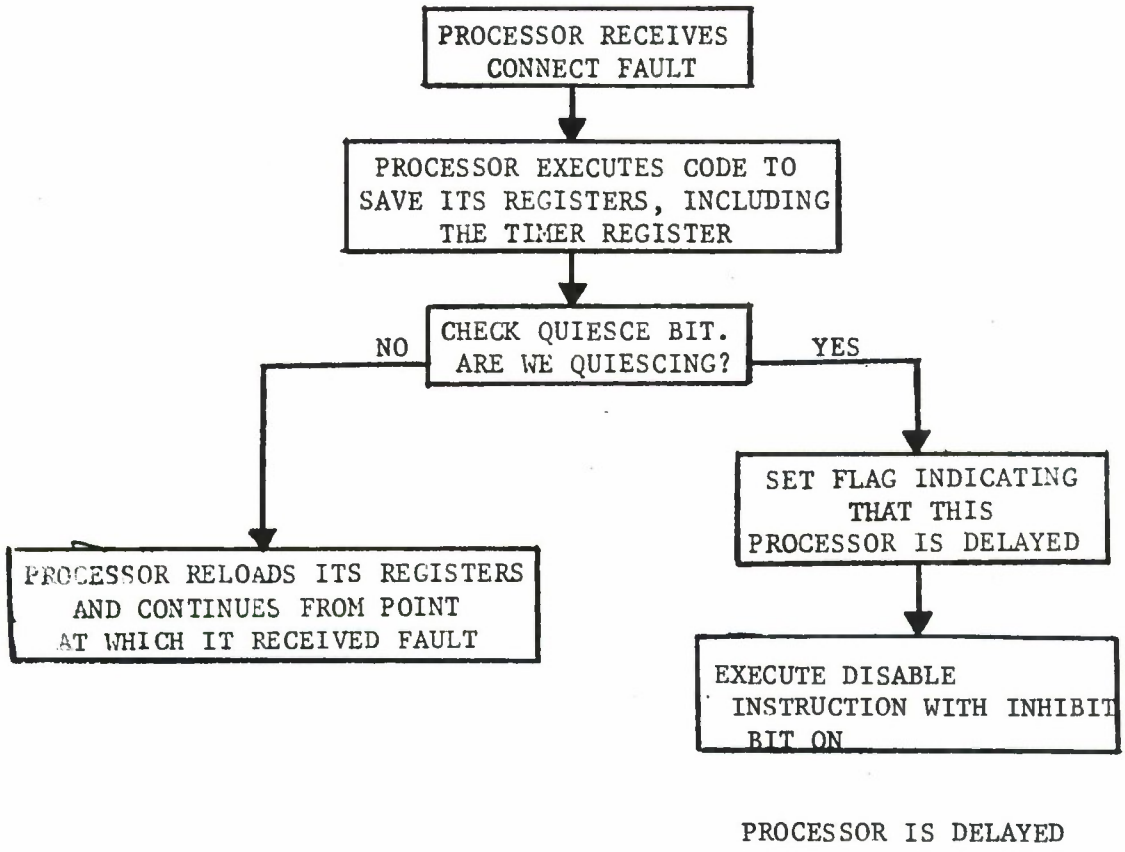

Figure II-5. Modifications to Connect Fault Routine for Processors Other Than Processor 0

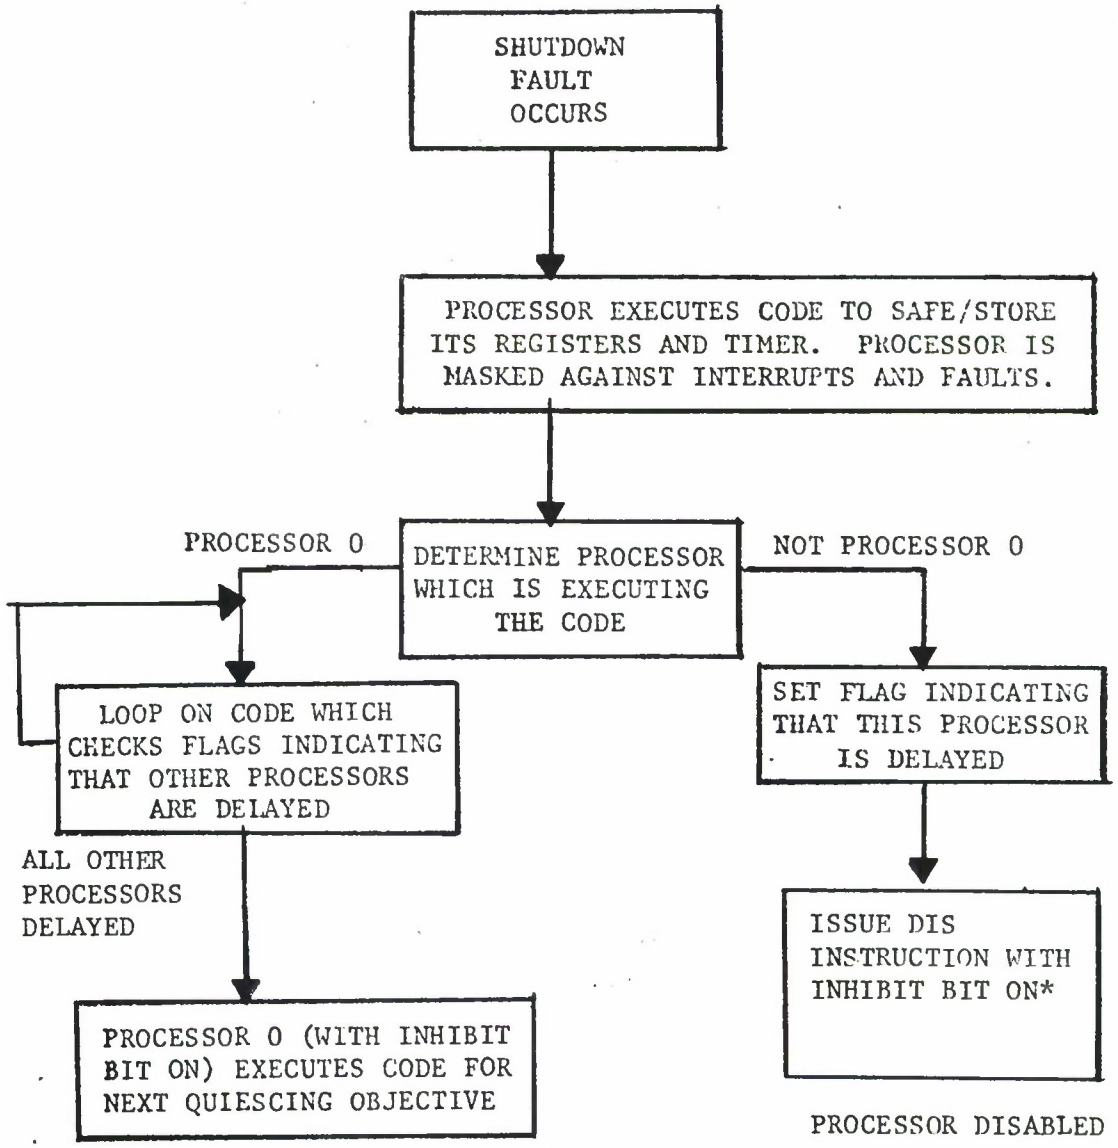

*By the time a processor executes the DIS instruction, the contents of the timer register have been saved but a timer runout fault could still be pending. If a processor responded to this fault during quiescing, the quiescing process could be aborted. Therefore, the precautionary measure of putting the inhibit bit on is recommended. This precautionary measure could be eliminated if a technique for canceling timer runout faults is developed.

Figure II-6. Processor Interpretation of Shutdown Fault as Quiesce Signal

logic of the shutdown fault routine within GCOS effectively ignores the occurrence of a shutdown fault. However, there is the chance that in future versions of GCOS the shutdown fault routine will be programmed to interpret the occurrence of a shutdown fault as being a meaningful signal (e.g., to clear cache memory).

Objective B

This objective is the stopping of all I/O operations between the IOM and H6000 memory. After objective A has been fulfilled all processors (except processor 0) are disabled. The IOM and system controller switches and registers are set such that only processor 0 responds to I/O interrupts. Processor 0 will execute the code that quiesces the I/O system.

The sequence of events that accompanies the completion of an I/O operation is as follows (see Figure II-7):

a. The peripheral controller signals the Input/Output Multiplexor (IOM) that a data transfer has completed.

b. The IOM stores status information concerning the I/O operation into the H6000 mailbox for the payload channel involved in the I/O operation. A mailbox is a set of contiguous memory locations within H6000 main memory. There are thirty-two mailboxes for each IOM - one mailbox for each payload channel. The IOM then does a read and clear of the memory location containing the interrupt multiplexor word (IMW) associated with this interrupt. There are four IMW words for each IOM. One IMW word is associated with each interrupt type. Thirty-two bits within an IMW word are meaningful-one bit for every payload channel. Next the IOM sets the IMW bit for the payload channel involved in the I/O operation. The IMW word is then stored back into H6000 memory.

c. The IOM sets the appropriate bit on in the execute interrupt register of the system controller. The execute interrupt register contains, at most, sixteen meaningful bits - one bit corresponding to each interrupt type on each IOM. The system controller has now been informed that an interrupt is requested.

d. The system controller checks the interrupt mask register. If this interrupt is not masked, then the system controller sends an interrupt present signal to Processor 0.

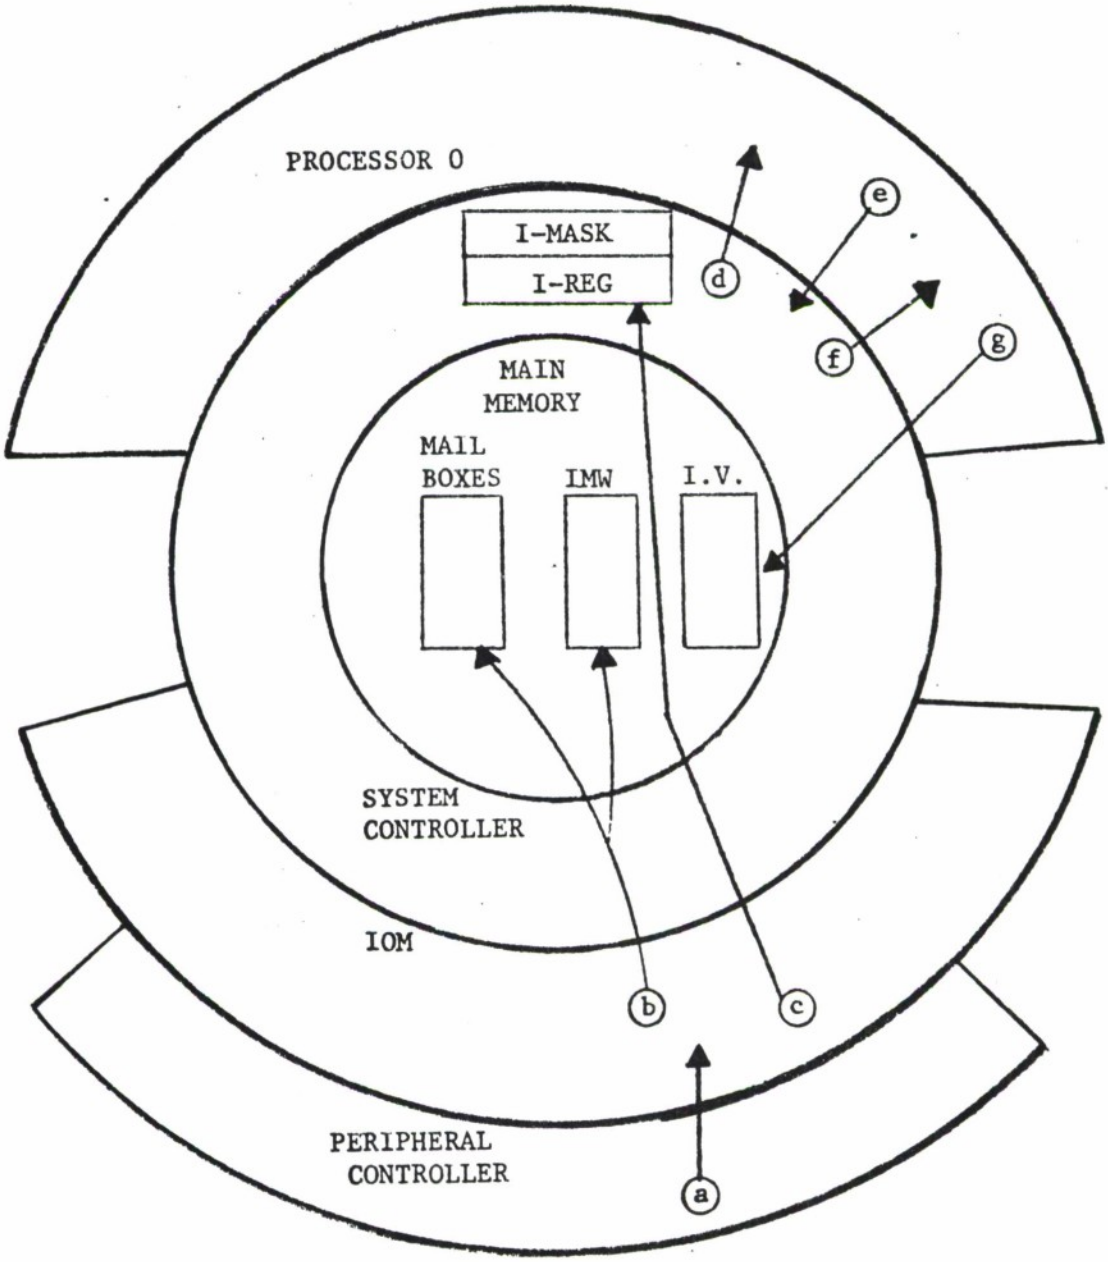

Figure II-7. H6000 I/O Sequence

e. Processor 0 sends an interrupt acknowledge signal back to the system controller as soon as this processor is ready to accept the interrupt.

f. The system controller passes the address of the interrupt vector entry (indexed by IOM number and interrupt type) for this interrupt to processor 0.

g. Processor 0 begins executing the instructions of the interrupt routine for this interrupt. The address of the routine is contained in the interrupt vector entry for this interrupt. The logic of the interrupt routine utilizes the contents of the appropriate IMW word (indexed by IOM number and interrupt type) to determine which payload channels have interrupts pending.

One method for fulfilling the objective of terminating all I/O operations consists of modifying a GCOS routine: the Input/Output Supervisor (IOS). A section of code within IOS, called STIO, initiates all I/O operations to peripherals attached to the IOM. An I/O quiesce flag could be added which would be checked by this section of code whenever it is called to initiate an I/O operation. When all outstanding I/O operations have been completed, then IOM activity has been quiesced (see Figure II-8).

A similar strategy could be devised for GCOS module M750 (the WWMCCS version of module DNET), which initiates I/O operations to the Datanet 355 by means of an IOM. However, the I/O initiation (CONNECT instruction) may not be centralized in module M750. There may be more than one section of code in M750 which initiates I/O operations to the DN-355. (In H6000 systems with extended memory, all connects to the DN-355 are made from code within STIO.)

There are two disadvantages to this first method:

a. It must be verified that there are no other sections of code within GCOS that initiate I/O operations. In the current version of GCOS there is only one section of code that initiates I/O operations, but future versions of GCOS may not adhere to this convention.

b. This alternative requires GCOS modification.

A second approach to quiescing input/output activity through

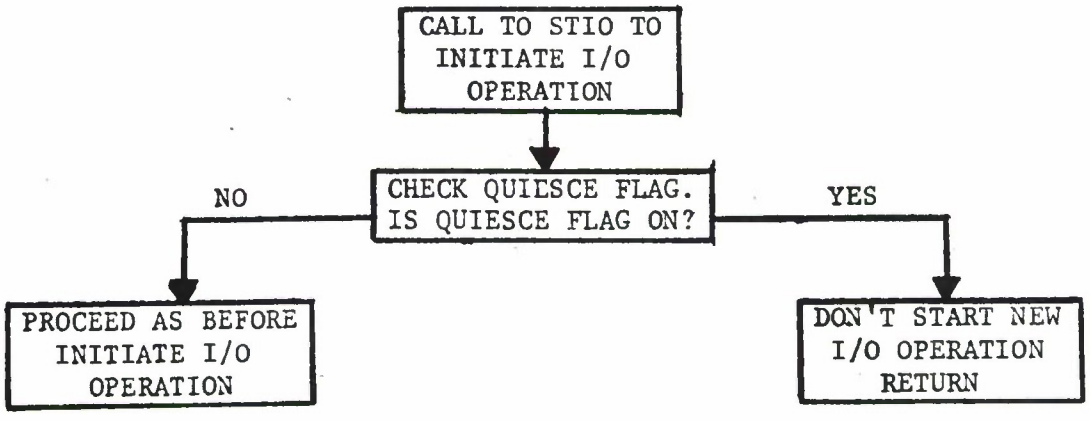

Figure II-8. Quiescing I/O at GCOS Routine-STIO

the IOM is for the H6000 quiescing routine, being executed by processor 0 with interrupts masked, to perform the following checks.

a. Determine which I/O operations are outstanding for each payload channel (i.e., a CIOC instruction has been executed) by following the chain of I/O data structures originating in the Primary and Secondary Configuration Tables.

b. Check the IMW words in the .CRIMW table*, in order to determine when the interrupt associated with the completion of each outstanding I/O operation is pending. Recall that upon the completion of a physical I/O operation, an IOM sets a bit in the Execute Interrupt Register of the system controller for the appropriate interrupt type. The IOM also sets a flag in the appropriate entry for this payload channel in table .CRIMW in 6000 memory. A pending I/O interrupt for an I/O operation means, therefore, that the data transfer to/from 6000 memory has finished, the I/O status information has been placed into 6000 memory, the appropriate bit in the Execute Interrupt Cell of the system controller has been set on, but processor 0 has not yet responded to this interrupt request by executing the instructions of the GCOS I/O interrupt routine.

Thus the quiescing routine can determine when the interrupts associated with all outstanding I/O operations are pending. This situation of pending interrupts can be recreated when this level is eventually restarted by means of the SMIC (set memory controller interrupt cells) instruction which is executed by 6000 processors.

---

* One exception to this procedure of determining that all I/O operations are in an interrupt pending state should be noted. It is possible to have several I/O operations outstanding on a single payload channel at one time (e.g., tape rewind operations). A special type interrupt will be associated with the completion of each of these I/O operations. In order to determine when all special interrupts are pending for this payload channel, it would be necessary to examine both the IMW words and the status control word pointer contained in the mailbox for this payload channel.

Although this method for quiescing 6000 IOM I/O activity has the advantage that only quiescing software is executed after the request for a level change has been made, there are certain disadvantages:

a. The algorithm for determining which I/O operations are outstanding is dependent upon the format of GCOS I/O data structures and the precision of information contained within them. Any change in those data structures could require a change in the algorithm.

b. The algorithm for determining from GCOS data structures all those I/O operations that are outstanding could become complicated, in that special cases for different types of I/O operations on various devices must be taken into account.

A third alternative for quiescing I/O activity would be for the 6000 quiesce software to wait long enough to guarantee that all I/O operations have finished and that the interrupt associated with each one of them is pending. Upon returning to this security level, this state of pending interrupts can be re-created by means of the SMIC instruction. The disadvantage of this alternative is in waiting too long; the quiescing software may continue to wait long after the interrupts associated with all outstanding I/O operations are pending. The time lost in waiting too long could be significant in an installation which utilizes tape drives, since a single tape I/O operation can last 7 minutes (e.g., forward file space involving a pass over the entire tape). Unnecessary waiting could be reduced by employing on a limited basis some of the tactics of the second alternative. For example, a check could be made to determine if any forward file space operations are outstanding. If none are outstanding, it should not be necessary to wait as long as 7 minutes. If a forward file space operation is outstanding, then the quiesce software could periodically check the appropriate IMW word in order to determine when the forward file space has completed.

The Honeywell Micro-Programmable Controllers (MPCs) for tapes and disks present some quiescing problems that warrant special consideration. MPC firmware can be categorized as:

1. Basic MPC firmware - general routines for communication with IOM, which are implemented in read-only memory.

2. Application dependent - for dealing with a specific type of device. These routines reside in (read/write) main memory.

The MPC routines will accept an instruction which puts the controller into a suspended mode. This instruction can come from H6000 software, and can be used to quiesce the MPC processor, before reading the MPC memory. Once all I/O between H6000 main memory and the IOM has ceased, the H6000 quiesce software (being executed by processor 0) can send the MPC a suspend controller command, wait until the MPC is in suspend mode, then read and save (if necessary) the contents of the MPC memory.

Quiescing the Datanet - 355

Quiescing the DN-355 presents some special problems, since the DN-355 (1) can operate independently of the H6000 computer and (2) is responsible for interfacing with users at terminals. There are a number of ways that the DN-355 can be quiesced. The method selected will depend upon the characteristics of a specific installation.

Method 1: Installation operating procedures are such that a warning can be given to terminal users in sufficient time so that they can save their files and log-off. In this case, it is reasonable to use the same DN-355 for more than one classification level of terminals. The DN-355 memory must be purged between levels by a certified DN-355 purge program. No modification to DN-355 software is necessary. Figure II-9 relates the quiescing of the DN-355 to the 6000 quiescing procedure. In this method all users should have logged-off before the 6000 quiesce software is executed. Any program for which a user hasn't disconnected himself is treated as if the DN-355 has gone down. This will mean, in most cases, that the program is aborted abnormally.

Method 2: In this situation it is assumed that there is not sufficient time to give terminal users a warning that quiescing is about to begin. It is necessary that 6000 quiescing begin immediately. This method consists of breaking the communication link between the 6000 and 355 and allowing them to quiesce independently. The 6000 quiescing will proceed as before and DN-355 software will be responsive to the terminal users. Modifications to DN-355 software are necessary. The primary advantage here is that the 6000 does not have to wait for terminal users to log-off before it begins to quiesce. Figure II-10 flow charts the logic for this method of quiescing the DN-355.

In this procedure the 6000 quiescing can begin before terminal users enter "save files" and "log-off" messages. These messages are retained in the DN-355 and are sent to the 6000 once control is returned again to this level. GCOS will then take the appropriate

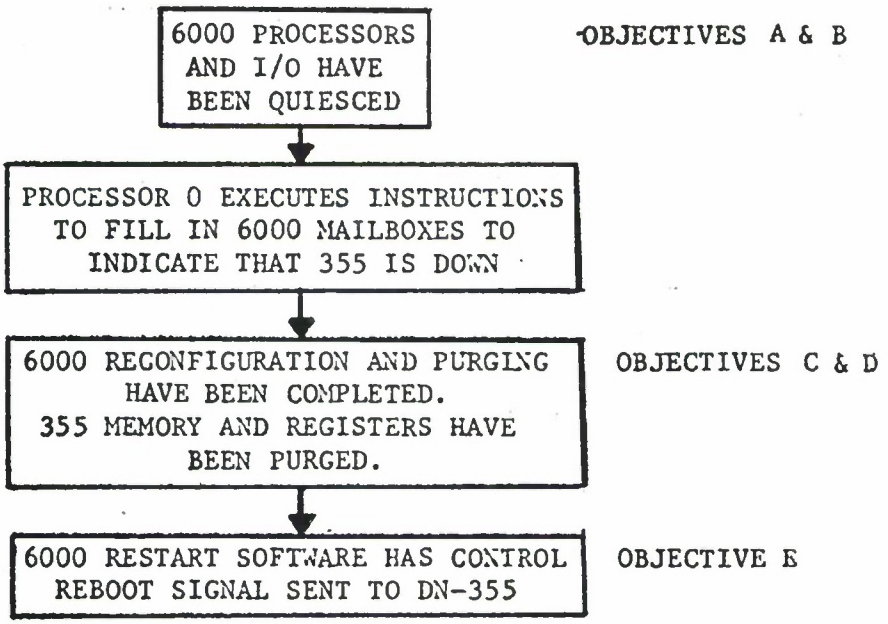

Figure II-9. Quiescing the DN-355
Method 1

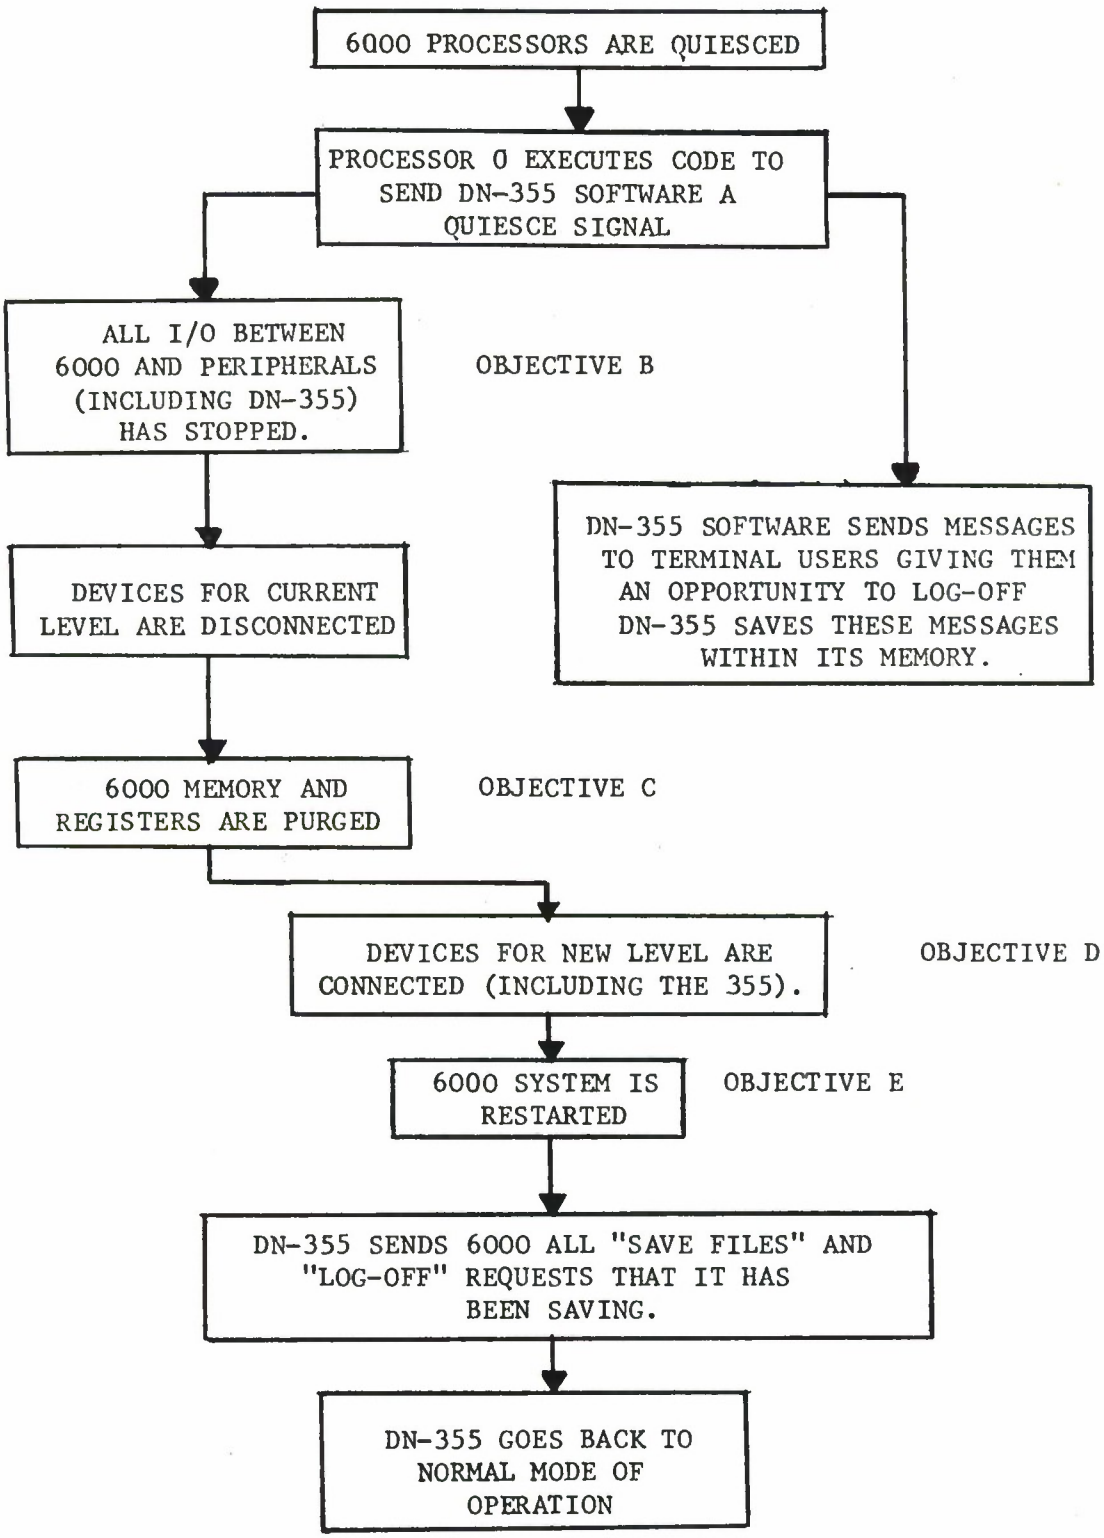

Figure II-10. Quiescing the DN-355
Method 2

action in response to these messages. A potential problem with this procedure is insufficient buffer space.

Objective C

Objective C is to perform any operations necessary for later device repositioning (i.e., repositioning peripheral devices to their position at the time of quiescing, when control is eventually returned to this level) and signaling the mini when the quiescing has been completed. Note that at this time all normal I/O activity has ceased and only processor 0 is active. Device repositioning operations may include: (a) remembering which disks are mounted on each spindle, and the current seek position; and (b) rewinding tapes to the beginning, and remembering which tapes are mounted on each drive and the current position of each tape. MPC main memory contains positional, status, and accounting data on devices. The contents of memory can be saved when quiescing and then reloaded when starting up; there are instructions which allow H6000 software to do this. (It may only be necessary to reload the original usage statistics for a level into the MPC.) The status and usage statistics for peripheral devices, along with the contents of H6000 memory, can be stored on a state-save file located on one of the disks for this security level. Figure II-11 contains a logical flowchart of the operations performed by processor 0.

If the controllers for peripheral devices are dedicated to one classification level, then some of the operations described to satisfy objective C are not necessary. It is not necessary to rewind tapes (this could save 2 minutes in the worst case) or to save the positional and status information within MPC main memory. The MPC's must be electronically disconnected from the H6000 IOM and it is still necessary to save the contents of H6000 main memory. Another possibility is to dedicate peripheral devices to a classification level and to share controllers between more than one level. In this situation tapes do not have to be rewound. The peripheral devices can be electronically disconnected from the MPC. However, the positional and status information within MPC's must be saved, since the MPC main memory will be purged.

Tapes: The reconfiguration of tape drives does present some problems. When quiescing a classification level, the H6000 quiesce software must rewind all tapes on shared tape drives and remember the original tape positions. This tape positional information is contained in the GCOS Secondary System Configuration Tables for tapes (file number, block number) and is, therefore, available to the H6000 quiesce software. It must be dynamically verified that

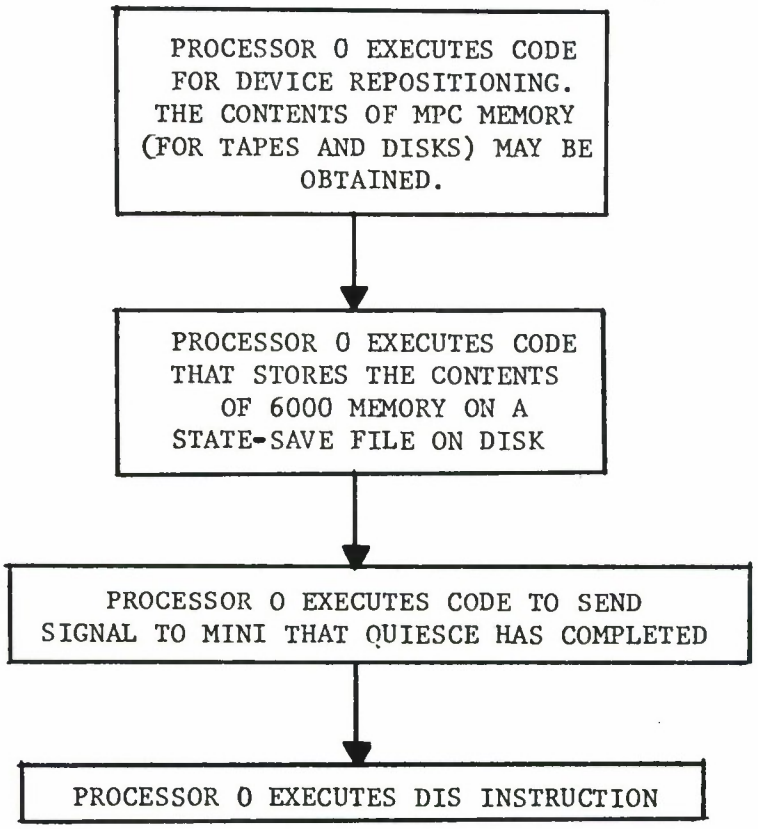

Figure II-11. Logic for Device Repositioning (Objective C)

device positional information is always accurate, even when the GCOS Exception Processing modules are being executed. When returning to this level, the H6000 restart software must return each tape to the position it occupied when the switch was made. As with disks it may be desirable to remount tapes on the same drives in order to eliminate modification of GCOS I/O queues and data structures. Micro-programmable controllers, called the Honeywell MTH 500, are supported by WWMCCS Release 5.0.

GCOS does have a facility for repositioning tapes to their original positions. The GCOS Exception Processing Routines perform the repositioning. When restarting a new classification level, these GEPR modules could be used to reposition those tapes on shared drives. This repositioning could proceed concurrently with other normal H6000 processing, after the H6000 restart routine has finished. Thus, the H6000 processing would not have to wait for the restart routine to finish positioning the tapes (in the worst case, this time delay could be 7 minutes). Only those slave programs that access tapes being repositioned must wait. Utilizing the GEPR modules to reposition the tape is, of course, a GCOS dependent activity.

The H6000 IOM configuration panel contains manually settable switches which define the source of the bootload program (tape or cards). These switches are:

a. Bootload source - tape or cards.

b. Tape channel number 0 if bootload source is tape.

c. Card channel number - if bootload source is cards.

d. Bootload port - defines port number of system controller through which connects are made to the IOM from a central processor.

Since the mechanisms already exist, it is recommended that the source for the restart and purge bootload programs be card decks or tape. If the mini is interfaced to the H6000 to appear as a tape drive, then the mini could serve as the source of the bootload program. Utilizing a disk as a bootload source could lessen manual intervention by eliminating the loading of card decks by the operator. However, this technique would require modification of existing bootload mechanisms.

Objective D

The actions required to fulfill objective D, which must be certified correct, consist of electronically disconnecting devices attached at the current classification level, clearing H6000 memory, registers and controllers, and connecting the devices for the new classification level (see Figure II-12). The MPC's and DN-355 present special problems with respect to purging. The H6000 software can send the MPC basic firmware a command to write into the MPC main (read/write) memory. Therefore, the H6000 purge program can assume responsibility for clearing the MPC memory. However, the MPC basic firmware must be certified, since this basic firmware does the actual storing into MPC main memory.

It has already been mentioned that MPC memory consists of two components - a read only component and a read/write component (which must be purged). Since it is theoretically possible to store information that could subvert H6000 purge activity, in the read/write component, the MPC read/write memory must be purged. Two potential solutions to this problem are possible. The threat could be met by having the minicomputer electronically disconnect the power supply to the MPC main memory. If this can be done without damaging any portion of the MPC or peripheral devices, then the read/write memory component, which is volatile MOS, would be purged. The second approach is to connect the minicomputer to the halt button on an MPC. Before the H6000 clear program is loaded, the minicomputer puts the MPC into a halt state (i.e., stops the MPC processor from executing instructions). The H6000 software can send commands to an MPC. These commands are interpreted and executed by firmware routines that reside in MPC read-only memory. One of these commands is a request to bootload a program from H6000 memory into MPC read/write memory. The H6000 clear program could therefore send the MPC (which has been halted) a bootload command, which would cause a certified MPC clear program to be loaded into MPC read/write memory and executed. A potential problem with the second approach is that the H6000 purge program and the MPC firmware are being trusted to correctly load and branch to the MPC purge program, and must therefore be certified correct.

The threat of improper information storage in the DN-355 memory must also be addressed. One possible solution is similar to the second approach discussed above for clearing MPC memory. The minicomputer is connected electronically to the halt switch of the DN-355. Before the H6000 purge program is loaded, the DN-355 is put into a halt state (i.e., the DN-355 central processor stops executing instructions). The H6000 purge program can then send a

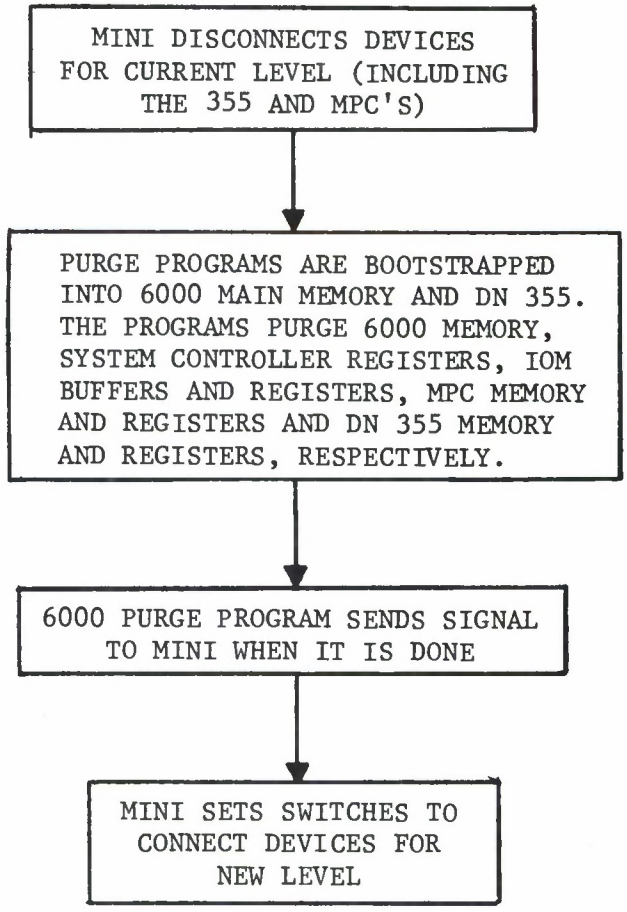

Figure II-12. Disconnecting Data Paths, Clearing Memory and Establishing New Security Level (Objective D)

bootload command to the DN-355 which would cause the following sequence of events to occur: (a) the starting component of a certified DN-355 bootload program is transferred into DN-355 memory from H6000 memory; (b) the DN-355 central processor is removed from the halt state; (c) the DN-355 central processor begins to execute the certified DN-355 bootload program which purges DN-355 memory and registers. It should be verified that the DN-355 CPU executes no other instructions between the instant at which it is removed from the halt state and the point at which it begins executing the DN-355 bootload program. It is also necessary that the operation of the direct interface adapter (DIA), the hardware interface between the H6000 and DN-355, be studied to insure that the proposed DN-355 clear procedure is correct. It may be necessary to halt the DIA before starting the DN-355 bootload procedure, since halting the DN-355 processor may not stop the DIA from continuing data transfer between H6000 and DN-355 memory. If MPC's or DN-355's can be dedicated to one classification level, then the necessity of purging them is, of course, eliminated.

Some peripheral devices may be dedicated to a single classification level. In this case, the device need only be electronically disconnected by the minicomputer, after all I/O operations between this device and H6000 memory have stopped. Other devices may be shared between two or more classification levels (see Figure II-13). The use of a device for more than one classification level requires the coordination of the computer operators, minicomputer, and H6000 quiescing/restarting software. The actions to be performed for each type of shared peripheral device are discussed next.

a. Card Reader: Computer operator must remove cards belonging to the old classification level before a switch is made to a new level. The cards must be repositioned in the reader exactly as before when returning to the new level. This should be a straightforward operation.

b. Printer: Computer operator must remove the paper belonging to the old classification level before the switch to the new one is made. When returning to the old level, the paper must be repositioned as before. The printer ribbon may also have to be changed.

c. Disk: For a shared spindle (drive), the operator must remove the pack belonging to the old classification level. When returning to a level, all packs which were mounted for that level at the time the switch occurred should be remounted. It is recommended that the packs be remounted

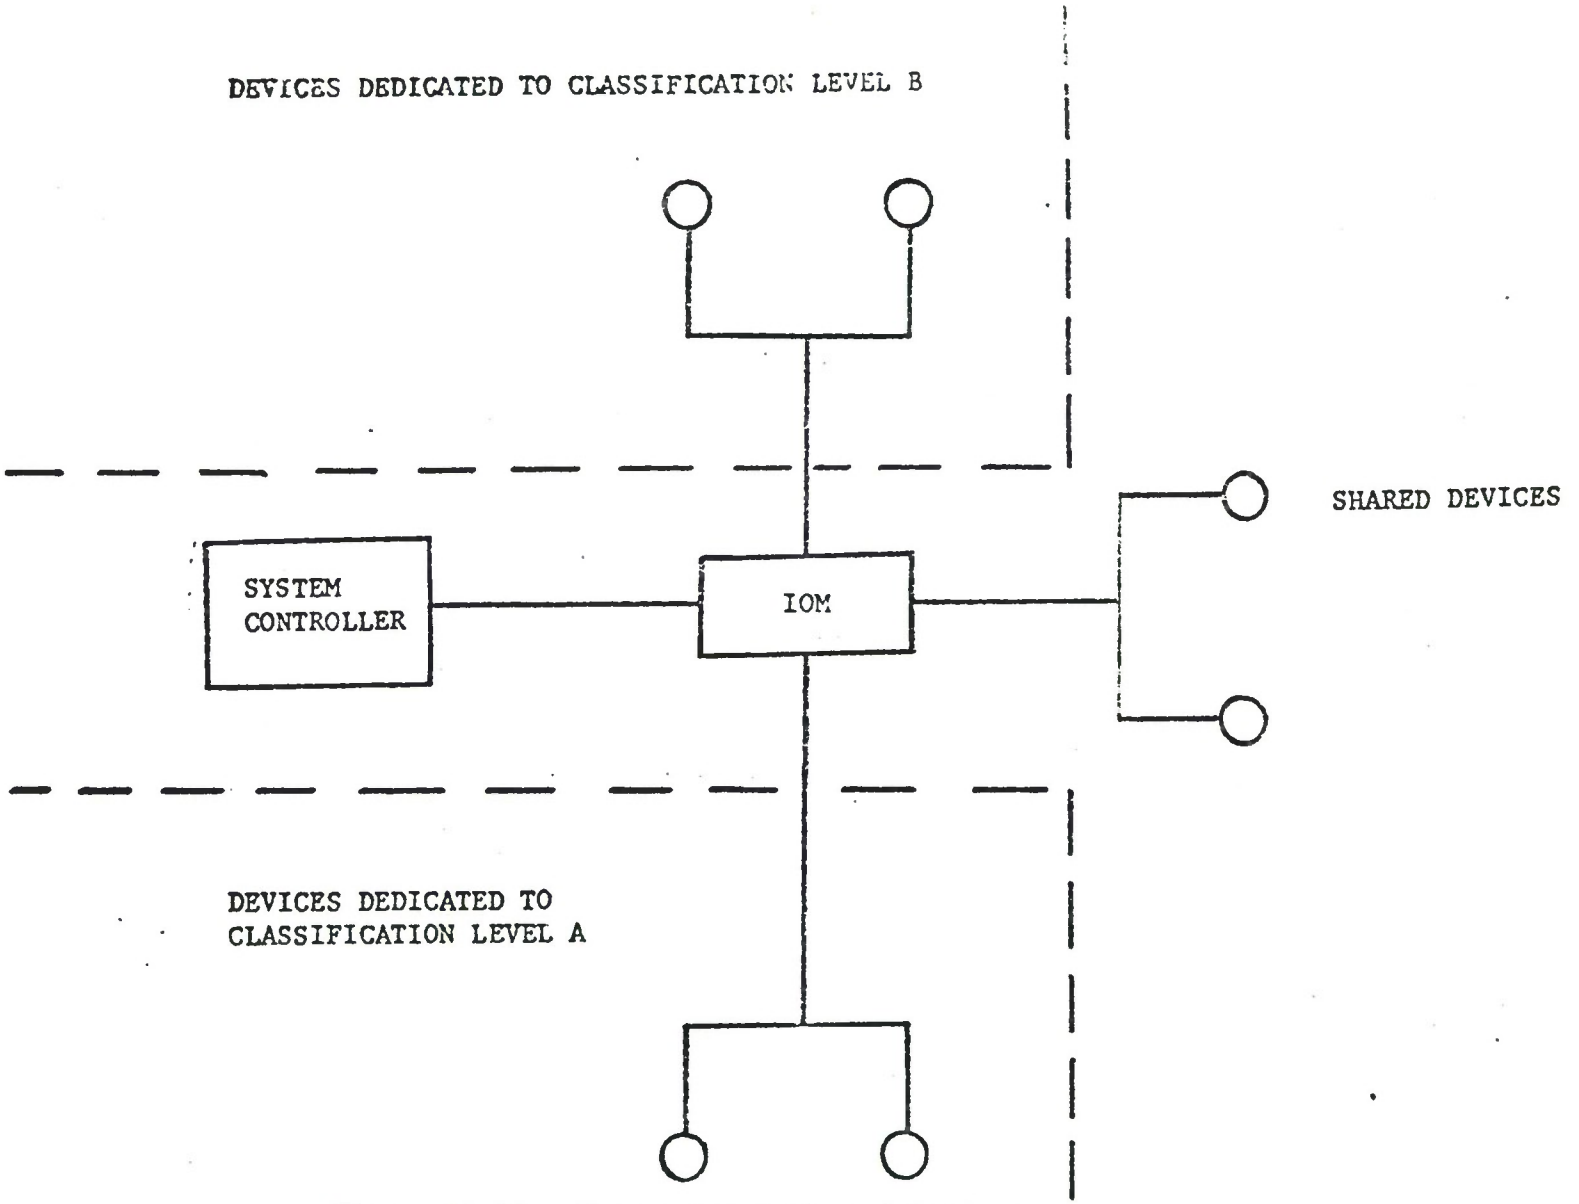

Figure II-13. Shared and Dedicated Devices

on the same spindles. The H6000 restart software should verify that the operator has correctly remounted the disks.

Objective E

Objective E is to restart the system so that processing at a new classification level can begin. It is intended that processing at the new classification level resume at the point at which it was interrupted (quiesced), as if the system had paused for only a moment. Time dependent operations within GCOS are based upon the elapsed time since system start-up, which is maintained in a GCOS data structure named .CRTOD. This data structure is updated whenever the interval timer register of a central processing unit is loaded. Within the context of the Jobstream Separator, the value within .CRTOD is interpreted to mean elapsed time since start-up of this level system. Since the value of .CRTOD and the value within the timer register are saved, the applications that are dependent upon .CRTOD should run as before. Any application programs that do not adhere to these GCOS conventions (i.e., programs that are dependent upon absolute time of day) may run incorrectly.

Once the restart bootload routine has been loaded, processor 0 begins execution. The shared MPC's are sent "BOOTLOAD CONTROL STORE" commands. When the MPC for the disk containing the system state-save file becomes operational, the H6000 restart software can read the data and instructions for the new level into H6000 memory. The shared devices, such as tapes and disks, are then repositioned. The I/O activity for the new level system is restarted. The actions performed here will depend upon the method used to satisfy objective B (i.e., by calls to GCOS routine STIO or by recreating the pending interrupts by means of the SMIC instruction). Processor 0 restarts the other processors so that they begin executing instructions at the point at which they responded to the shutdown request. Processor 0 reloads its registers and begins executing instructions at the point where it responded to the shutdown request. The new level has now been restarted (see Figure II-14).

APPROACH II: QUIESCING AT THE INTERFACE BETWEEN GCOS HARD CORE MONITOR AND SLAVE PROGRAMS

GCOS module MPOPM is a privileged slave that allocates memory resource to slave programs and communicates with the computer operator. Quiescing at the interface between the Hard Core Monitor and slave programs would involve changing MPOPM; one technique for

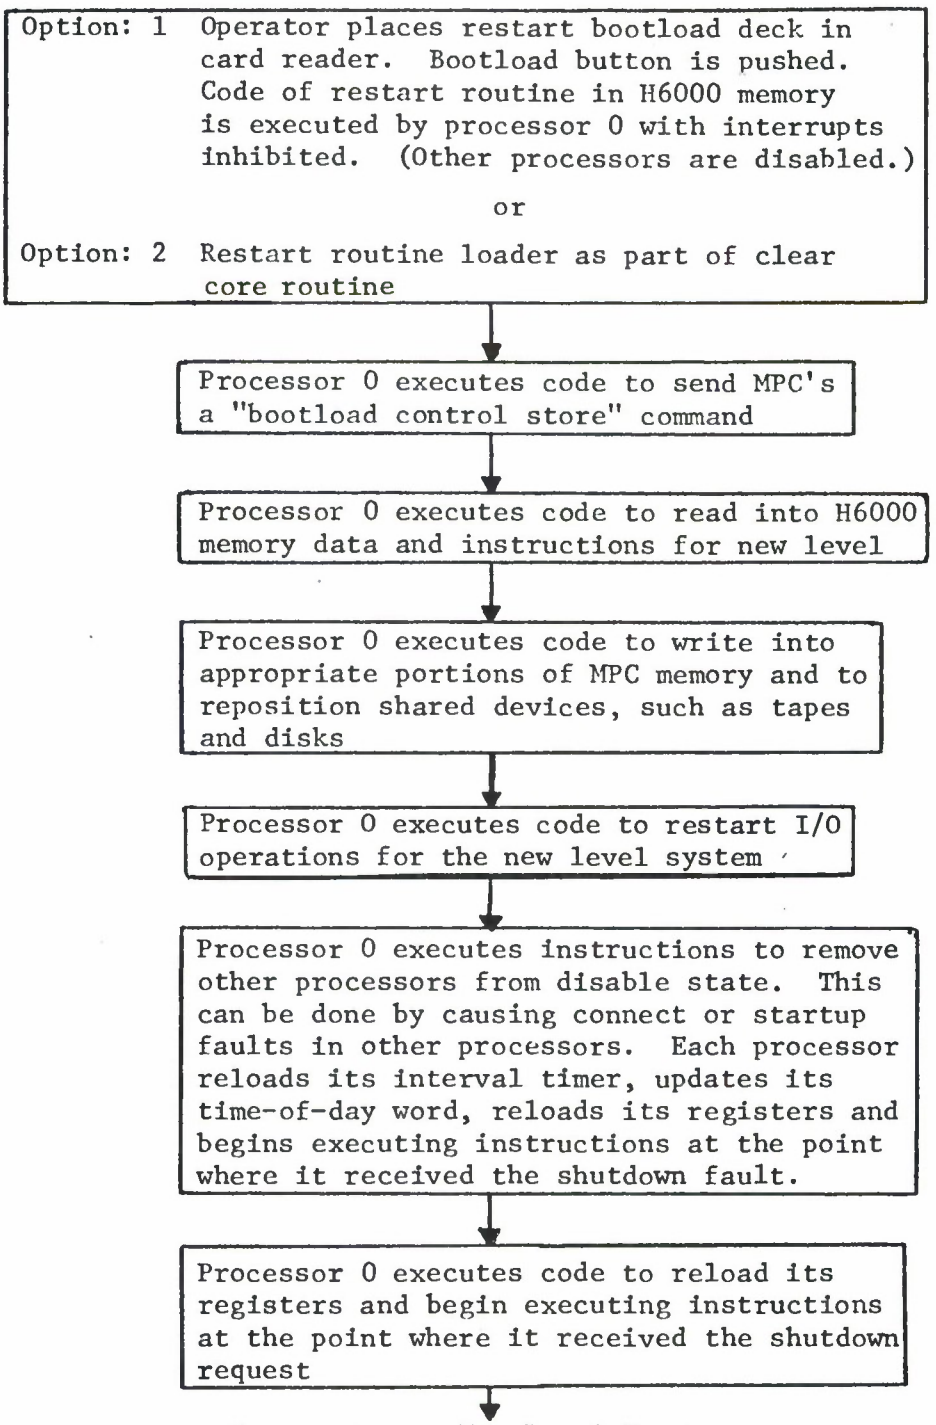

Figure II-14. Restart Logic (Objective E)

quiescing at this level is discussed here. MPOPM would be modified to accept a new console verb, i.e., a new operator command. Upon receiving this command (QUIESCE), MPOPM would enter quiesce mode. MPOPM would insure that all slave programs (including privileged slaves) are swapped out. Once this has been done, the only slave program that can be executed by a processor is MPOPM, so that the only I/O operations that can be initiated are those from MPOPM. MPOPM would recognize this situation and send a signal to the mini indicating that the H6000 system has been quiesced (see Figure II-15).

Objectives B, C, and D, would be done in the same manner as if quiescing took place at the interface between the H6000 hardware and the hard core monitor. With respect to the MPC's, disk repositioning would not be necessary, since I/O requests to disk would not be terminated after the head positioning and before the data transfer. The accounting data in the MPC must be saved and restored. The clearing of MPC memory and registers would be the same as before. Figure II-16 indicates the logic for fulfilling objective E.

There are several outstanding problems with this approach, which should be mentioned.

a. Some privileged slaves, such as Time-Sharing and GEOT, cannot be easily swapped. Some modifications must be made to facilitate the swapping of these slaves within some "reasonable" time period.

b. A technique for quiescing the exception processing modules (GEPR) must be developed.

c. Privileged slaves in master mode cannot be swapped, according to the current policy. However, with the extended memory feature, the software can be changed to allow for the swapping of privileged slaves that are being executed in master mode. Future versions of GCOS may allow for this.

d. An upper time bound for quiescing, clearing and restarting utilizing this method cannot be given now. However, several observations can be made to establish a lower bound. The time bound for objectives A and B would be at least as large as that indicated in Table II-I (i.e., 7 minutes). The time bound for objectives C and E should be the same as that contained in Table II-I (i.e., 2.5 minutes). The time required for objective E would be about the same as before.

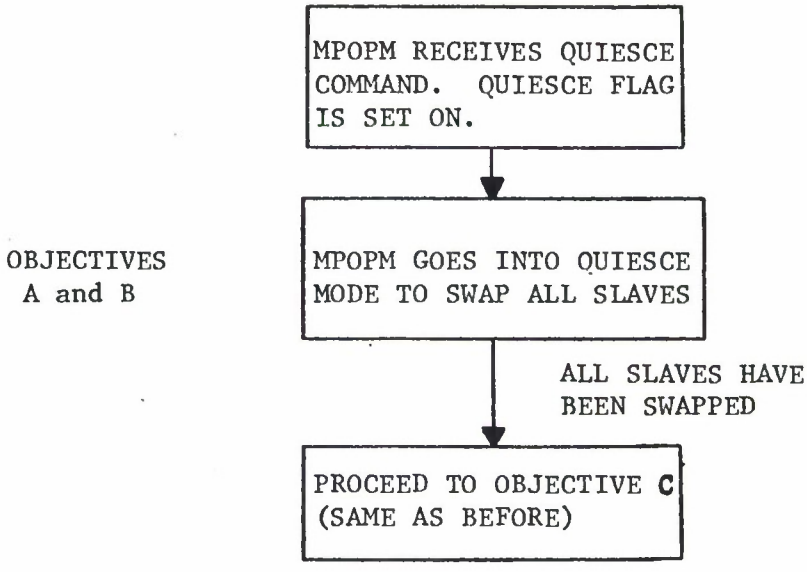

Figure II-15. MPOPM Fulfilling Objectives A and B

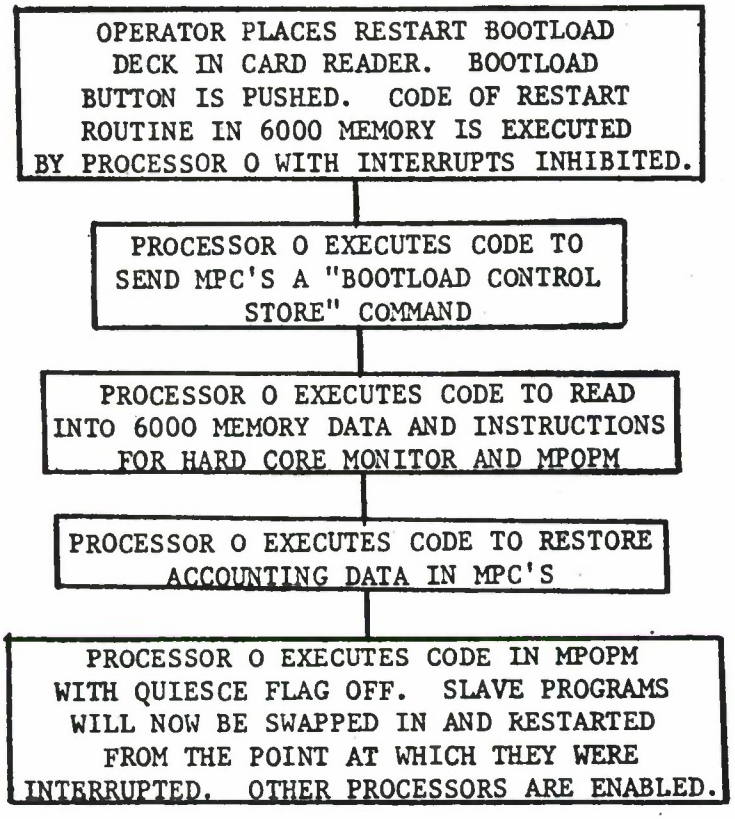

Figure II-16. Logic for Fulfilling Objective E

COMPARISON OF TWO APPROACHES

It is estimated that the software necessary to implement either approach would require approximately 18 man-months effort. The DN-355 software could require the most extensive modifications. This would occur if the policy of having DN-355 software retain terminal user's "save-files" and "log-off" messages is adopted. There are several differences between the two approaches that merit consideration. For the following discussion Approach I refers to the hardware-HCM interface and Approach II refers to the HCM-slave interface.

a. The time required to quiesce, purge, and restart will be less with Approach I. It is not yet known if the time difference between the two approaches will be significant.

b. The interface at the hardware level is better-defined and should change less frequently. The interface at the HCM-slave level is subject to programmer style and creativity. The interface is not as precise and conventions may be disregarded. Thus, Approach II would be more difficult to maintain.

c. Approach II requires more GCOS modifications. Most of these modifications would be to MPOPM and related modules.

d. In implementing Approach II, the GCOS facilities that already exist for quiescing slave programs and debugging could be utilized. Honeywell personnel indicated that a commercial release of GCOS will contain many of the necessary facilities for quiescing programs. These facilities may be available in future WWMCCS GCOS releases.

e. More Honeywell expertise is available at the HCM-slave interface. Thus, the implementers of Approach II could expect to obtain help and advice more easily.

APPENDIX III

ACCESS CONTROL SWITCHES

The JSS system employs an array of switches interposed between the minicomputer and various system components, storage devices and/or their controllers. This appendix describes a design for that array of switches, and outlines some alternatives which were considered in its development.

Each switch is actually composed of two distinct parts: an access control switch driver (ACSD), which is one of many drivers physically resident on a single interface board, and an access control switch (ACS) proper, which is connected to the device or its controller. The ACSDs, and through them, the ACSs, are under the minicomputer's program control -- they are set on and off by a minicomputer task that generates a digital signal corresponding to the appropriate action. The minicomputer also has the ability to sense the status of each ACS individually. In fact, the mini performs periodic comparisons of actual ACS settings with desired ACSD settings. If a mismatch is detected, the offending ACS is automatically located and the appropriate error-handling action is taken.

ELECTRONIC VERSUS ELECTROMECHANICAL SWITCH

Two types of switching devices can be used to implement the ACS: electronic (e.g., TTL circuitry) or electromechanical (e.g., relays).

Relative to electromechanical switches, electronic switches operate faster, require less power, are easier to design and have a higher degree of reliability. Electromechanical switches are more costly (approximately $5 per switch and $50 per power supply). Electronic switches are clearly favored.

ACCESS CONTROL SWITCH DRIVER/MONITOR

The access control switch driver/monitor is shown in Figure III-1. Although the general purpose interface board is shown separate from the mini, it would typically be incorporated into the mini's mainframe. The components shown in the diagram are the mini, the board, one of many ACSDs on board, the ACS under the control of an ACSD, and a typical device (disk or tape).

The box labelled "typical device" does not actually show how or where the ACS is connected to the device. Although Figure III-1 shows the ACS separate from the device, the ACS will be attached and reside inside the device's housing. The device is temporarily disconnected by controlling either the ON-LINE/OFF-LINE lines or by accessing the READ/WRITE INHIBIT lines. The ON-LINE/OFF-LINE method is preferred because it is easier to access, involves less logic circuitry and provides a sure means of totally disabling the device.

A blowup of the ACSD is shown in Figure III-1 to permit a high-level inspection of its internal logic. Power sources, relays, and logic subcomponents are intentionally not included. A brief description of the internal control lines follows.

The SENSE line is used to check if the signals from the minicomputer have reached the target access control switch and have been obeyed. Faulty circuitry or wiring can be detected by the minicomputer by comparing the INHIBIT and SENSE signals; the ENABLE signal is also used in the comparison -- it should be off when the DISABLE line is on. If the result of the comparison indicates a fault detection, the ALARM line is activated and an interrupt is sent to the JSS minicomputer to indicate that an error condition has arisen and appropriate action must be taken (e.g., ring a bell or light a display lamp). The MONITOR line, which connects the ACSD to a storage device, is used to ensure that no I/O operations are required of a device that has supposedly been disabled. If a device's INHIBIT line is on and I/O activity is noted (i.e., obviously the ACS switch has failed and has permitted I/O operations to occur on a drive not intended for use at the present time), the error condition is stored in one of the interface board's status registers and an interrupt is sent to the minicomputer for further attention. The INTERLOCK line insures the proper physical connection between the ACSD and the ACS. If the connection between the two is not complete, then the switch and/or its connections are faulty and the alarm function is activated.

ACS-DEVICE INTERFACE

An ACS is connected to the JSS minicomputer via a General Purpose Interface Board. Typically, these boards can handle 48 external lines and can therefore address and control 48 ACSs. The minicomputer maintains an internal table for each ACS: the address of the ACS, the address of its ACSD, and the physical address and mnemonic of the target device.

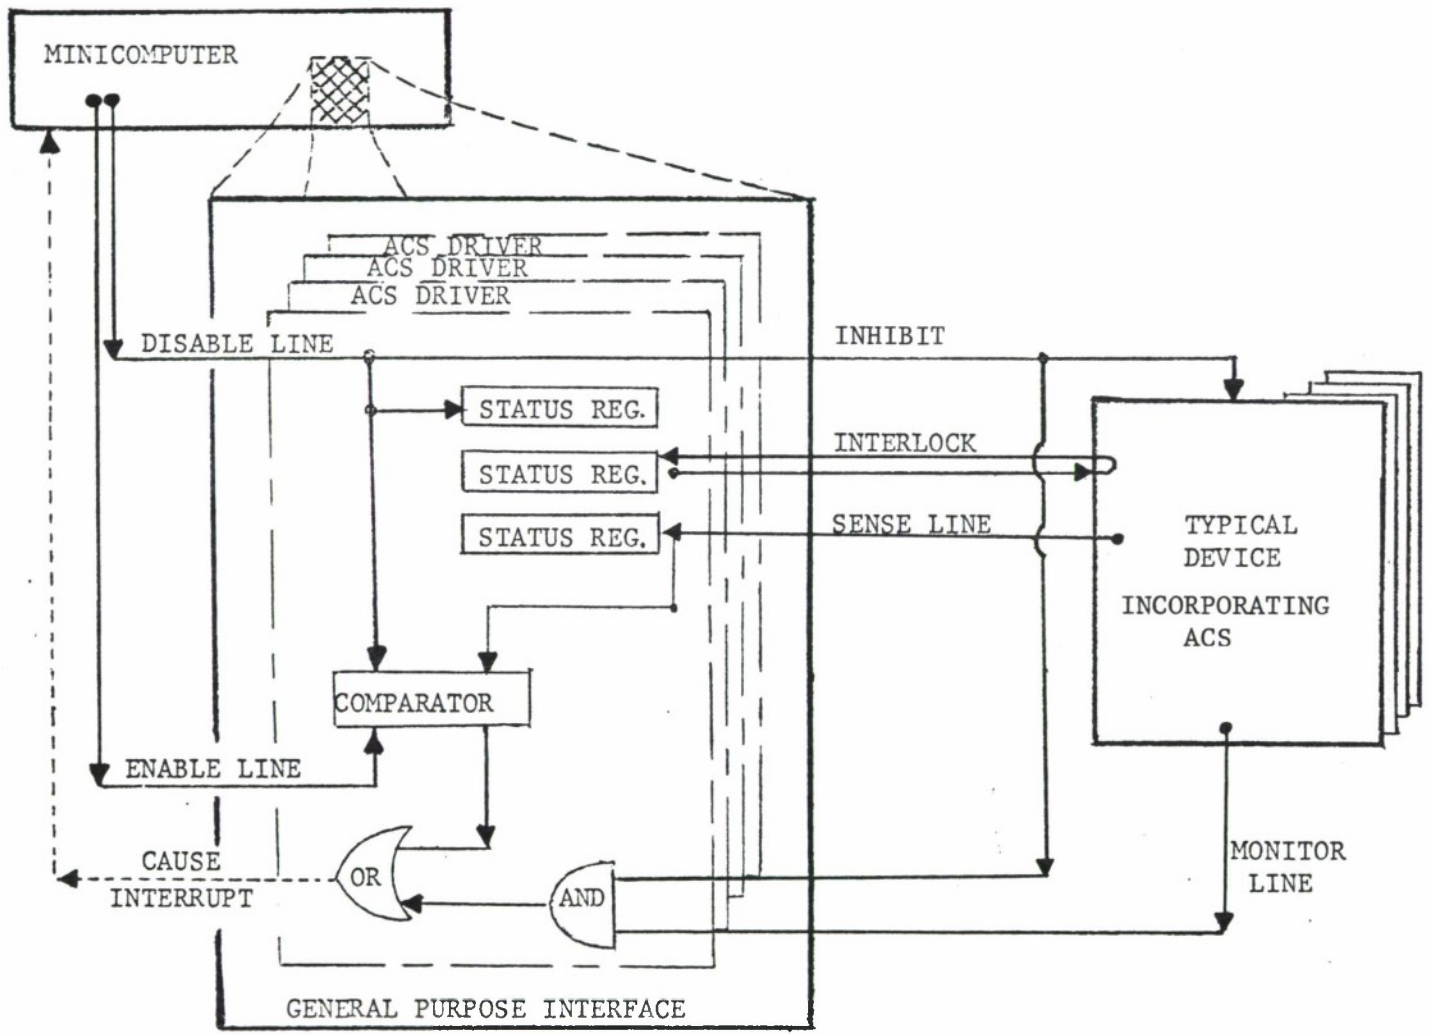

Figure III-1. Access Control Switch Driver/Monitor

There are two interfaces that must be considered: the ACSD-minicomputer interface and the ACS-device interface. The first is a rather straightforward engineering task common to most real-time or process control minicomputer applications, and will not be treated here. The second interface, the interconnection of the ACS to either a device or its controller, requires additional engineering and hardware, and, depending upon the point of interface chosen, affects the total cost of a JSS system.

Three alternative ACS-device connections are possible, as indicated in Figure III-2, which shows a schematic of a typical peripheral device to CPU connection via a device controller and line adapter. (Assume for the remainder of this discussion that only disk drives and disk controllers are under consideration.)

The first connection possibility is to install an ACS at Point A, inside the device controller at a point where system information from all devices is available at a common node. The advantage of this technique is that the number of ACS devices and the associated circuitry and cabling to disable the device are minimized (1 per controller). The installation costs would be small and all circuitry is localized at the controller level. The disadvantages of this method are obvious and numerous, but the approach is still a viable solution. First, Honeywell would have to concur in (minor) modification of the internal circuitry of the controller in order to allow an external enable/disable control signal to be interfaced with Honeywell control lines, in addition to incorporating a mechanism for monitoring the status of each control signal. Second, as the ACS becomes more complicated, it becomes more costly to design, build, and install. Third, one must also consider the possibility that the internal design of the device controller may change with new releases of Honeywell hardware, thereby negating ACS interconnections. Lastly, physical size constraints must be accounted for if the ACS is to be installed within the controller.

The second connection alternative illustrated in Figure III-2 is to connect the ACS at Point B, between the line adapters and the device controller. This method would necessitate using twice as many ACS boards as the first scheme (Point A), since most device controllers permit a maximum of two line controller connections. The advantages and disadvantages of this technique are similar to those of the first alternative. The major differences are that this method requires more installation connections and the ACS circuitry is less complicated.

The third connection alternative shown in Figure III-2 is to connect the ACS directly to each drive at Point C. The major disadvantages of this method are that one ACS is needed for each device,

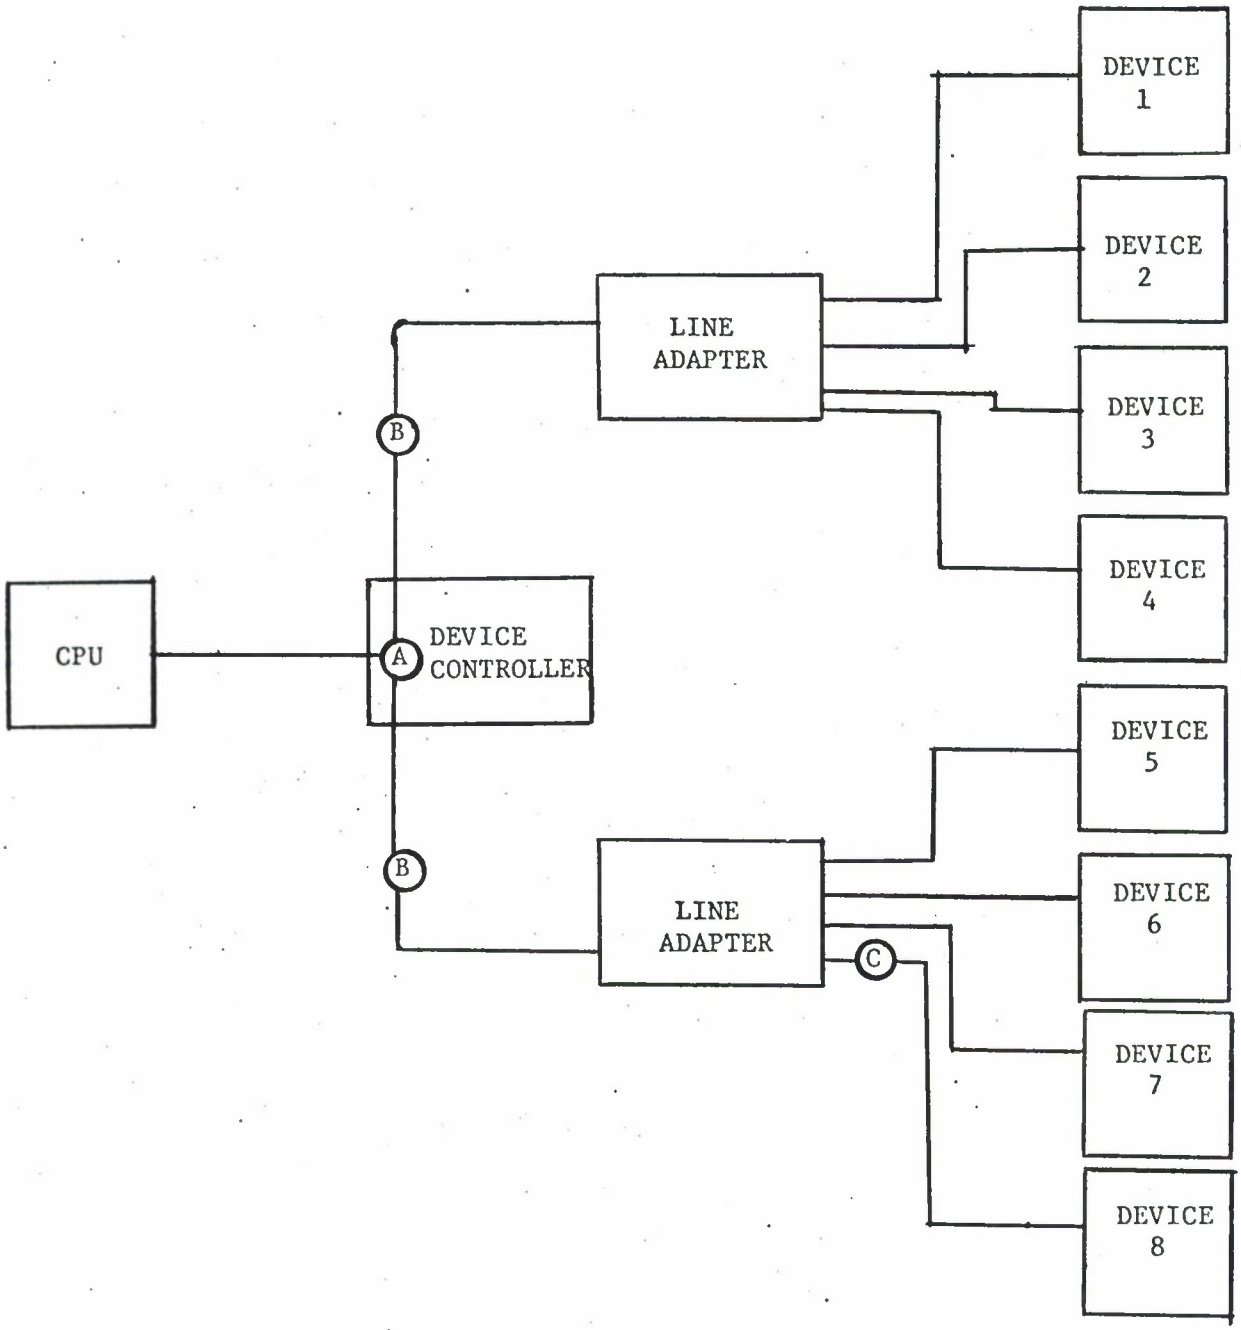

Figure III-2. Typical CPU to Peripheral Architecture

and that device circuitry must be modified and/or accessed, which increases the installation cost. The advantages, however, are that the internal activation of an ACS within a device makes JSS control conceptually more absolute, and that the ACS circuitry is very straightforward and could be designed, built, and tested very quickly.

Of the three interconnection alternatives posed, it is recommended that the third approach, device-level interface at Point C, be taken. This approach requires access to internals, however the relatively simple "clip on" ACS circuit should not impact disk or tape drive performance in any way. This method also allows the system configuration to change without necessitating extensive modification to the ACSD or the ACS circuitry, as might be the case if the controller housed the switches. Should a device be added to the configuration, it would only be necessary to add an ACS to the device and connect it to the appropriate line on the mini.

Although connecting the ACS in the controller at Point A would seem to be the most centralized location, the complexity of the logic design necessary to accomplish the interface would not compare favorably with the other alternatives. In addition, practical considerations dictate that Honeywell will most probably offer stiff opposition to this method -- they would deem it a costly intrusion into their proprietary hardware and would not accept responsibility for possible malfunctions that stem from JSS-ACS interface. (This opinion was expressed by several members of the Honeywell staff.) The second approach, line adapter interface at Point B, is less complicated (in terms of design and installation) than the controller interface (at Point A) but is more complicated than switch implementation at the device level. Lastly, since the first two interface approaches require access to centralized control lines, an access switch failure would compromise more drives than would the Point C alternative, where an ACS failure would affect a single drive.

APPENDIX IV

SYSTEM CONTROL/DISPLAY PANEL

This appendix describes the function of the System Status Display/Control Panel. The panel serves two primary functions. First, it allows the operator to input control information concerning the dynamic reconfiguration of system peripherals during color-change activity. Second, it provides a visual status summary report of all system devices during both color-change and normal activity, and in addition, summarizes the progress of JSS color-change activity. The control and display functions of the panel are treated separately in the ensuing subsections.

DISPLAY FUNCTION

Two types of data concerning system peripherals are of interest to the computer operator. One type is the status of the peripherals during system operation; the other is their status during a color-change. The status panel described below has been designed to show the first type of data, and to summarize the second.

The operator does not require the detailed progress of the color-change to be displayed on the panel. The operator only needs to know: a color change has been initiated, the mini is responding to the color-change request, the resulting security level of the system, and the final status and disposition of the peripheral devices. If the operator needs detailed information, it is displayed to him on the JSS's console. An elaborate panel need not be designed to encode this information. With the advent of a JSS controlled color-change, such a panel would more likely seem to be a confusing array of flashing lights than a helpful display.

A status panel is shown in Figure IV-1. This panel is intended to show the operator the following information:

System status - down, off, on, security level, color-change occurring

Peripherals - down, off, on, security level, color-change occurring

Disks/tapes - mounted, security level

At the top of the panel is a light showing the overall system status. This enables the operator to tell at a glance whether or

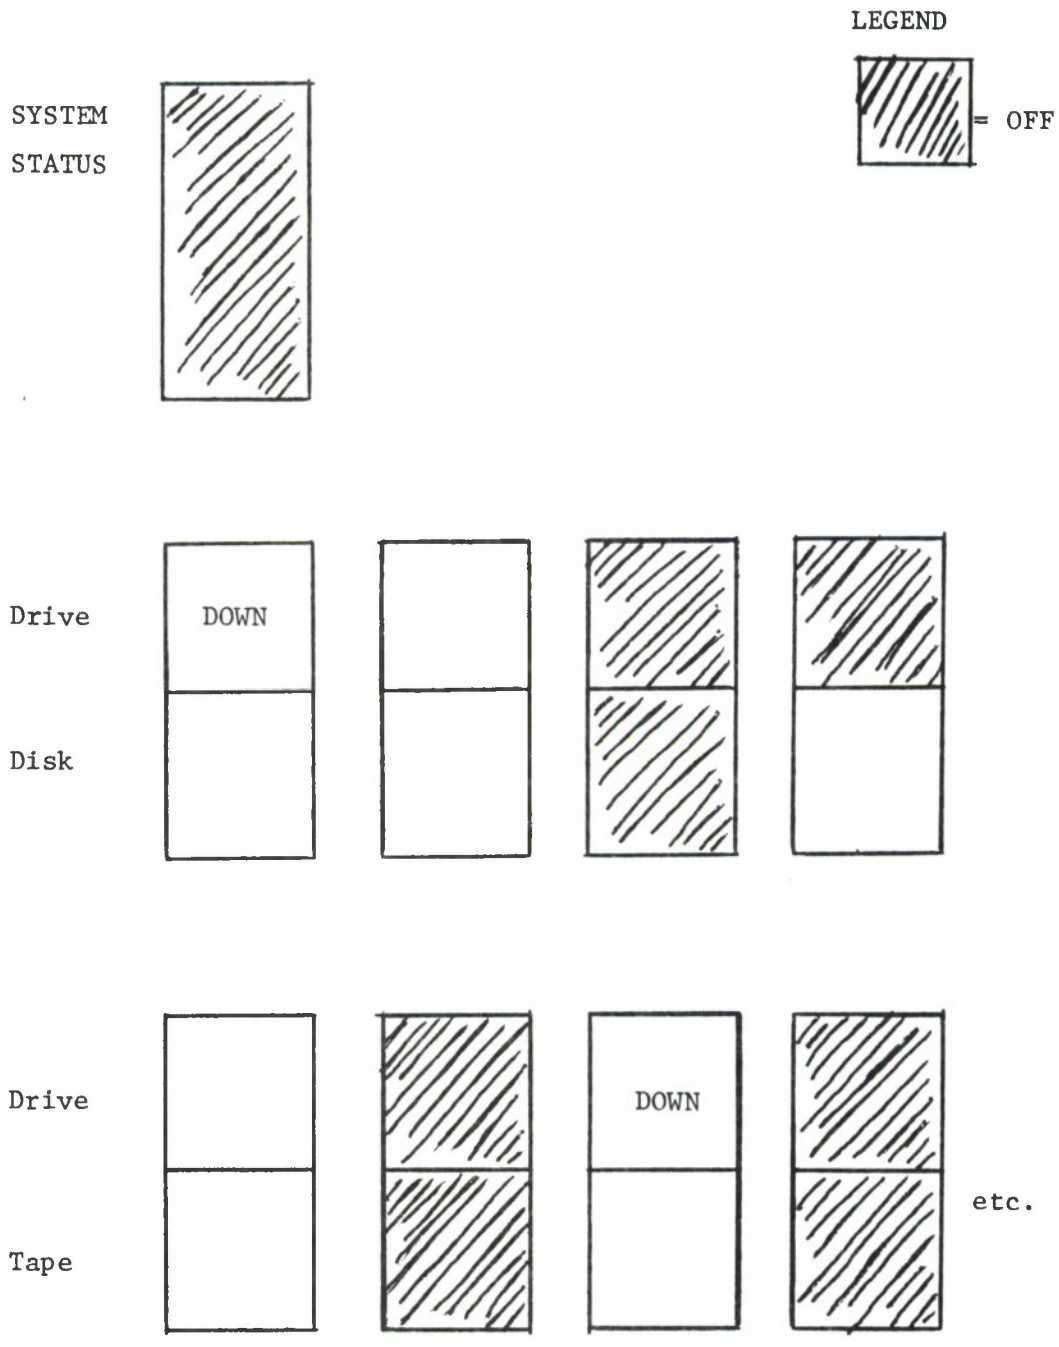

Figure IV-1. Sample Peripheral Status Panel

not the system is up, and if so, the current security level. These are the codes used to convey this information:

System Status	Light Status
Down	Off, with DOWN displayed
Off	Off
On - Unclassified	On, white
On - Secret	On, yellow
On - Top Secret	On, red
Color-change in progress	On, blinking

The system status light blinks when a color-change is effected. The color of the light would not change until the facility's security level has changed and the color-change is complete.

Each peripheral device has a light associated with it coded in the same manner as the system panel. A peripheral device that is in a state of transition (i.e., being mounted/dismounted) would be indicated by a blinking light.

The peripheral-lights are numbered to correspond to numbers marked on the peripherals themselves. The number should be prominently displayed on each peripheral device, so the operator can see it without difficulty. The lights are arranged on the panel in rows which do not necessarily correspond to the physical layout of the facility. Thus, the room layout can be changed without causing the panel to become out-of-date and confusing. The lights are large enough for the operator to see when he is across the room; if he is unsure of the status of a peripheral, he does not have to leave it to check the display.

Associated with the light for each drive, there is a light for the tape or disk which can be mounted on that drive. If none is mounted, the light is not lit. If a disk/tape is mounted, the light is lit with the color appropriate to its security level. No blinking code is used for disks/tapes during a color-change, since they do not themselves change in status. It may be possible to have a disk mounted when the drive is not active; in this case the drive light would be off and the disk light on.

This panel can be expanded, using the same basic coding system, if additional peripherals are to be included in the facility.

The information conveyed by this panel could be shown via a tabular display on a CRT, or typed on a teletype. Since it is summary data, a panel of colored lights is recommended. The operator can tell at a glance which peripherals are in use, which are available, and whether or not a color-change is complete. He cannot tell the specific identification of a disk/tape on a particular drive. If these details are needed, the operator should be able to query the system using a CRT or teletype.

Supplementary Tabular Display

An operator requiring detailed information about the system status should be given the ability to query a file containing the system status, the identification of each peripheral and its status, the identification for tapes/disks mounted and their status, etc. The result of his query would be a display or list containing the relevant information. A sample of such a display is shown in Figure IV-2. Obviously, more information can be included if necessary.

If a CRT or similar display is available, the entire table can be shown at once. If the operator interface to the system is a teletype, he should have a choice of requesting the data for a subset of peripherals of particular interest to him, or the entire display.

Comparing Figures IV-1 and IV-2, one can more easily see the status of the system with the summary status display. The detailed table requires more careful attention, but provides useful information as well as serving as a backup to the status panel.

CONTROL FUNCTION

The JSS controls dynamic reconfiguration of system devices via the electronic access control switches connected from the control minicomputer to those devices. The JSS allows two convenient forms of reconfiguration control input: 1) the operator may enter a series of commands at the JSS console in order to specify the devices in question and the operations to be performed, or 2) the operator may manually set a toggle switch (or push button) on the Panel that corresponds to the desired device and operation to be performed. In either case, it is the JSS minicomputer which receives the input signal from either the JSS minicomputer console or status panel, decodes and ensures validity (i.e., timeliness, accuracy, etc.) and initiates the certified code to perform the

SYSTEM STATUS:

OPERATIONAL - SECRET

Disk Drives:		Disks:	
ID	STATUS	ID	STATUS
DØ1	DOWN	-	-
DØ2	OFF		-
DØ3	SECRET	DØ7	SECRET
DØ4	SECRET	-	-

Tape Drives:		Tapes:	
ID	STATUS	ID	STATUS
TØ5	OFF	-	-
TØ6	SECRET	TØ3	SECRET
TØ7	DOWN	-	-
TØ8	SECRET	TØ4	SECRET

Figure IV-2. Supplementary Tabular Display

desired action. The panel is merely a fast manual analog of the JSS command console -- the setting of a toggle switch simulates the entry of typed commands.

As can be seen in Figure IV-3, the panel does not have direct access via electrical connections to the Access Control Switches at each device. The JSS minicomputer "reads" the current settings of the switches by performing sensing and comparison operations and "writes" the status information to the control panel by selecting the appropriate set of display lights.

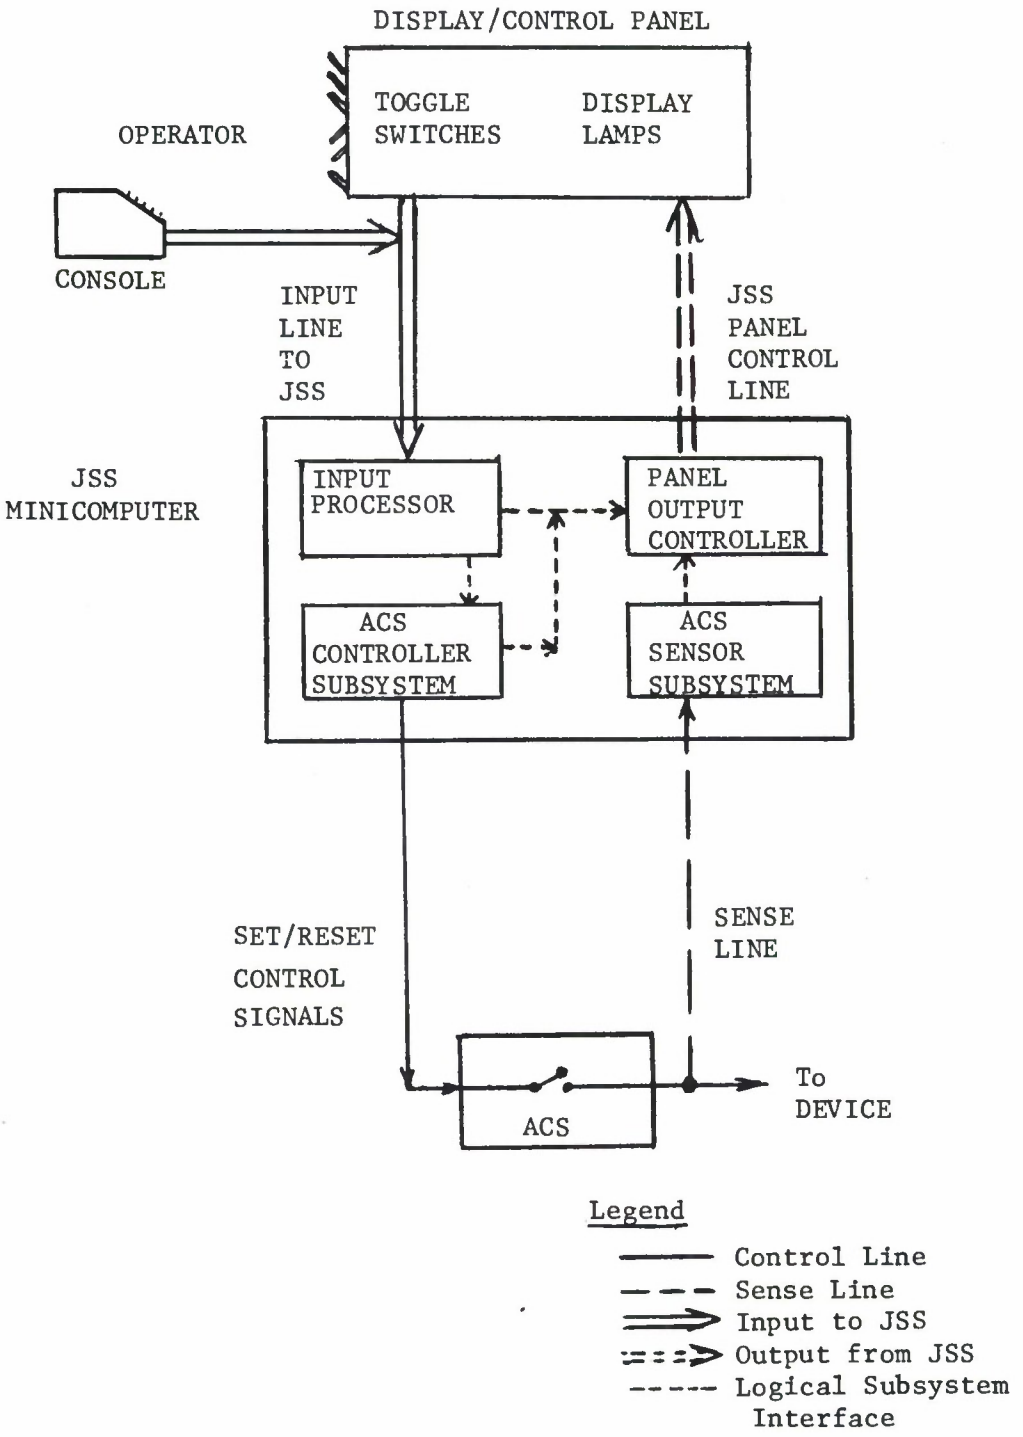

Figure IV-3. JSS Minicomputer - System Display/Control Panel Interaction

APPENDIX V

MACHINE-READABLE LABEL PROCESSING

Since the JSS involves dynamic reconfiguration of peripheral devices (i.e., disk and tape drives), a device may, during the course of daily operation, contain storage media of different security levels (e.g., a disk drive designated TOP SECRET an hour ago might be used as a SECRET drive after a color change). Although the JSS assures that peripheral devices are connected to the computer at the proper time and at the proper security level, it does not assure that the proper tape or disk is mounted on the correct drive at the correct time.

In order to reduce the possibility of operator error, in the mount and dismount process during normal processing periods in addition to the color change reconfiguration phase of the color change, it is proposed that a facility based on "machine readable label technology" be incorporated into the JSS system. This facility would serve as a means of monitoring the security level of all storage media handled by the WWMCCS computer, and would provide extensive control and verification of what is mounted where and when. Machine readable identification data would be attached to each tape and disk drive and to each tape reel and disk pack. The operator would be required to input the identification data to the jobstream separator, at mount time, by the simple, rapid, and error-free means of stroking a wand across the machine readable data.

CONCEPT OF OPERATION

The facility entails attaching adhesive machine readable labels to:

a. all storage media (tape reels and disk packs) to indicate the sensitivity of the data contained within. The identifier could be attached to either the protective case or to the handle or reel proper. The identifier could also include owner identification, or similar information that coult be used by a tape/disk library or other control points,

b. all storage devices (tapes or disk drives) to indicate the device class and device name.

An additional form of machine readable label (e.g., menu cards), either attached to each drive or carried by the operator, could be used as a convenient means of notifying the mini of completed actions or entering status information pertaining to each device. The cards might contain a list of human readable words (e.g., MOUNTED, DISMOUNTED), beside which are corresponding machine readable information that would be input by wanding the label.

To explain the use of machine readable labels and the associated hardware and software mechanisms, let us assume that a disk pack of level A is to be mounted on a particular disk drive that is currently unoccupied. The operator retrieves the desired pack and mounts it on the designated spindle. He then passes a "light wand" over the machine readable label attached to the pack. The wand "senses" the data on the label and stores it in a temporary buffer within the wand device. When the buffer is full (the label has been read), the JSS minicomputer system is signalled via an interrupt generated by the light wand device. A task in the mini then decodes the data on the label -- the mini now "knows" the security level of the pack being loaded. If the label data are illegible or the characters read indicate an invalid security level, a warning light is flashed on the system display panel to indicate the operation was unsuccessful. A successful label read operation is indicated by a "beep" at the wand outlet located near the operator. Assuming the pack label was read successfully, the operator next strokes the label attached to the drive now containing the mounted pack. The minicomputer, upon receiving another interrupt, decodes the label data from the drive and uses the encoded drive name to search an internal system table, the Environment Control Table (ECT). The mini checks whether the drive is one of those designated to be mounted/dismounted. If it is not, an error message is typed at the main console, the status of the system display is modified, and no "beep" is issued to the operator. Once the mini establishes that the correct drive has been chosen, it matches the known security level of the pack with the intended security level of the drive. If a match occurs, a successful mount operation has been performed. If the storage medium's security level does not match that of the drive, the mini can lock out the drive by disabling the "on-line" control line, signal an error condition to the operator at the main console, and modify the status of the system display panel.

As an additional security precaution to protect against the operator performing the "dual-stroke" operation (i.e., moving the wand across labels on the pack and drive) and then placing the pack on a different drive, the time duration between stroking the pack and the drive can be monitored by the minicomputer. Should the time-lag between strokes exceed some preset maximum (i.e., 15 seconds),

a warning message to the operator's console or reconfiguration panel would be issued.

DESCRIPTION OF TECHNOLOGY

Technology to support the above concept is readily available and is being widely applied commercially in point-of-sale applications. The machine readable data may be encoded on a magnetic stripe, or as bar codes or printed characters suitable for optical recognition. Other forms are also available, but the above three are the primary commercial contenders. Because magnetic stripe would offer problems in attaching and reading on other magnetic material or on metal, optical bar codes and optical character recognition are the most likely candidates.

Machine readable bar codes or characters can be imprinted on adhesive label material. The labels could be bordered in colors suitable for color-coded security designations. Human readable translations can be printed under bar codes, and full alphanumeric data can be encoded. Optically readable characters could be printed in OCR-font that is also human-readable, but to date only numerics and a few alphanumerics and special symbols can be read by wanding. Capturing the machine readable data in either case is a simple process of stroking a pencil-like wand across the label surface. Internal error checking and the addition of external check digits can make this reading process virtually error-free.

COMPONENTS

The principal components for a system based on either bar codes or optically readable characters include the following:

- Labels. Barring extreme requirements, label material should cost a few cents each.
- Label Printers. Labels can be generated under direct computer control or by human operator keyboard action. Printers would cost from three to five thousand dollars each, with the price heavily dependent upon the total number purchased at one time.
- Wand and Associated Electronics. A buffer, error checking logic, and control logic is usually a part of each wand. Each wand and its electronics will cost approximately a few hundred dollars, dependent upon the total number purchased.

- Communications. Communications between the wand and the jobstream separator minicomputer can be by hardwired cable. Cable costs depend on the number of wands and physical layout of a particular facility.

- Minicomputer Interface. A hardware interface for connecting wands with the minicomputer will be required.

- Minicomputer Software. Software for wand device handling should be minimal. Application software for verifying and controlling operator actions will depend on the extent of the procedures implemented.

BENEFITS

A facility for processing machine readable labels:

a. increases the mini's security related role during reconfiguration periods,

b. lessens the number of manual security related steps in the operator's checklist,

c. aids human factor considerations dealing with operator performance,

d. is based on reliable, efficient, and widely used technology (exceptionally high reliability rate for light wands, etc.), and

e. allows for fast and easy communication between operator and minicomputer control system.

TRADEOFF ANALYSIS

The use of machine readable labels as a vehicle for insuring proper mount/dismount action is optional; however, the requirements met by such a facility are not. If this feature is not incorporated in the JSS, some other mechanism, such as direct keyboard entry by the operator into the mini's console, will be required to take its place.

The most obvious "disadvantage" associated with this process is that it increases the cost of the basic JSS system. The bulk of the cost rests in the label-producing machine, whose basic price is

approximately $5000. This cost, however, represents an initial outlay (which can be discounted considerably depending upon the number of machines purchased) that can quickly be recouped because of the multiplicity of applications for machine readable labels (i.e., data management, tape/disk libraries, etc.). In fact, a single label producing machine at a WWMCCS site could serve not only all the needs at that site, but could service all WWMCCS requests for label generation. Regardless of the cost of the production machine, the price of an individual label (approximately one cent) is insignificant.

Additional implications and ramifications involving this facility are:

a. additional software development costs for control programs running in the mini that monitor the interrupts generated by reading labels. The development cost for this system may be less than the cost needed to develop an alternative solution (i.e., operator entry into mini's console) that would lack the ease and efficiency associated with the proposed system and would introduce more time-consuming security related steps;

b. additional hardware features -- wands, buffer, sockets, and interface boards. Total cost of hardware per wand-unit is approximately $500. The interface board, which costs $600, is interconnected to the minicomputer and can service several wands. Power for the wands is usually obtained from the mini's power supply;

c. increased memory required to handle label processing mechanism. A current system running on a Data General Nova, requires approximately 2K words of memory;

d. additional time required to connect hardware components to drives, mini, interface board, etc.;

e. a label identification code must be devised;

f. all drives and storage media must be labelled.

ADDITIONAL CONSIDERATIONS

The following is a partial list of alternatives, additional functions, and problem areas concerning the JSS label reading mechanism.

a. What happens when a "foreign" disk pack (or tape) is to be processed by JSS-controlled system? How is new label (possibly, a temporary one) produced? The mini operator should have the option to override the security level indicated in order to accept mislabelled or unlabelled tapes or packs.

b. Should a time limit be imposed between wanding operations? The JSS could monitor the time taken by the operator to perform the mounting, wanding, and start-drive operations. Should a preset maximum time period be exceeded, the mini could then signal the operator that the mount/dismount request has been aborted. The operator would then have to repeat the entire wanding sequence within the prescribed time period.

c. Should the label-reading mechanism not function properly, the operator must be given the capability of overriding the entire label checking process. The mini should maintain awareness of the status of this checking mechanism.

d. The question of how to handle multiple simultaneous mounts/dismounts of packs on various drives must be taken into account. If only one operator is mounting disks then there is no need to "lockout" one or more wands in the label reading process. However, if two or more operators are mounting and wanding disk packs then the mini's label reading task must be cognizant of the origin of the data (i.e., by prefacing the data taken from the pack label or drive label with a "bit sequence" indicating which wand is being used).

APPENDIX VI

SUPPORTING DATA FOR JSS COST ANALYSIS

This appendix serves as a supplement to the economic analysis presented in Volume I. The first part of the appendix describes the quantitative assumptions used to calculate the time (in minutes) lost per year by color changing under present conditions. The second part of the appendix calculates the money saved (net gain) over a three year period of JSS operation.

EXPECTED TIME SAVED BY JSS OPERATION

In order to calculate the expected time saved by JSS operation certain assumptions must be made.

Assumptions

- Current color change operations require 45 minutes -- this estimate is based on figures obtained from MAC and SAC. It includes a 30 minute period of system dead-time and 15 minutes for pre-color change preparation and post-color change restart. Although a 60 minute estimate may be more accurate, 45 minutes was chosen to account for the variability of the pre- and post- color change periods.

- 10-minute JSS color change duration -- this estimate eliminates the extremes of JSS operation (e.g., a 7-minute tape I/O operation or system reconfiguration involving the dismounting of all removable storage media). This figure is intentionally set at 10-minutes to compensate for any unforeseen delay factors; the expected duration is 5 minutes.

- An hour of WWMCCS system time costs $430 -- this estimate is based on data collected from CONAD (which has two H6060 systems valued at $429 per hour) and Air Training Command (which has one H6060 system valued at $398 per hour). Other WWMCCS sites, especially those that employ the larger H6080 systems (e.g., MAC and SAC) were unable to provide an hourly estimate. Typically, commercial rates for mid and full-sized Honeywell 6000 series systems or their equivalents (e.g., IBM, CDC, and UNIVAC systems) are in the range of

$400 to $750 per hour. For example, Decision Computing, a Cambridge, Massachusetts firm that houses an H6050 system, values an hour of system time at $600. Although the $430 estimate is used in this analysis as a cost guideline, since it was obtained from a WWMCCS site, it should be noted that any estimates in the $400 to $600 range could be justified, especially when dealing with the larger H6080 configurations that require more physical space, power, and manpower.

- Three color changes are performed per day -- this estimate is based on actual data collected at MAC and SAC. It also can be viewed as an extension of mission requirements at other AF WWMCCS sites in the near future.

- A 24 hour - 7 day week is assumed for WWMCCS system operation.

- The cost of the JSS system is $1.225 million -- this is the overall program cost for 6 WWMCCS sites over a 3 year operational period.

Calculations

For one site for one year; making

3	color changes per day, for
365	days a year, yields
1095	color changes per year at
45	minutes per current color change represents a total of
49275	minutes spent per year as opposed to
10	minutes per change with a JSS, which yields
10950	minutes per year -- a savings of
38325	minutes, or
638.75	hours per year. Since an hour of WWMCCS time costs
$430	the delay incurred represents a savings of
$274,662	per year per site.

For 6 sites over a three-year operational period:

638.75 hours saved at one site, for 1 year or

3832.5 hours saved at 6 sites, for 1 year, or

11497.5 hours for 3 years; a savings of

$4,943,925 for 3 years at 6 sites.

APPENDIX VII

DATA COLLECTION

The Jobstream Separator feasibility study and requirements analysis required data collection from various WWMCCS sites. This section describes the data collection methods employed and summarizes the information obtained.

The intent of the data collection effort was to gather information pertaining to the operational policies, workload statistics, and security requirements of AFWWMCCS sites. The data collected was intended to aid the requirements analysis in three ways:

1. To help determine the functional and operational capabilities required of a Jobstream Separator System at each site.

2. To generate a composite (overall) view of the Jobstream Separator System as common requirements were enumerated.

3. To aid in the preparation of a cost benefit analysis associated with the introduction of a Jobstream Separator System at AFWWMCCS sites.

This section presents the data collected and focuses on the following topics concerning each site:

1. System configuration.

2. Mission requirements -- present and future.

3. Special hardware features.

4. Future growth patterns -- introduction of new equipment or planned changes in mission requirements.

5. Operational problems presented to WWMCCS ADPE by the JSS.

6. Opinions concerning need for improved secure processing and personal thoughts of operators and managers regarding JSS-operator interaction and JSS functional requirements.

DATA COLLECTION METHODS

The methods used for data collection were:

1. Visits to WWMCCS sites.

2. Telephone calls.

3. AUTODIN messages.

Questionnaire

A questionnaire was designed prior to the first site visit in order to formalize and structure the numerous questions that were to be asked. In the preliminary stages of data collection, the questionnaire served more as a guideline for the interviewer than as a "fill-in" for the interviewee. The form enabled the interviewer to ask questions and formulate new ones as needed. After each trip, however, it was found that a more structured format was needed in order to economize interview time, make the questionnaire more tractable, and account for any unexpected developments or discoveries that might occur during a questioning period. Thus, the questionnaire went through several design iterations based on newly collected data, and as more experience was accrued, the questionnaire became more polished and comprehensive.

Telephone Survey

Whenever possible, MITRE and ESD representatives made personal visits to the WWMCCS sites. However, after having visited three large AFWWMCCS sites (AFDSC, MAC, and SAC), an additional site survey of a typical "small" WWMCCS installation (AFDSDC) was conducted via telephone conversation. The questionnaire was used as a guideline for the questioning.

AUTODIN Messages

Other AFWWMCCS sites were also surveyed by use of a short form of the questionnaire. The questions were:

1. Describe your WWMCCS system configuration (e.g., number of disks, tapes, system controllers, IOMs, MPC's, DN-355's).

2. Briefly describe your site's mission. Are you a fully operational site?

3. Briefly describe types of application/development programs run.

4. Do you handle multilevel information (i.e., Unclassified, Secret, Top Secret) or are you dedicated to a single level?

5. Do you "color-change" (i.e., change from one security level to another)? If so, how many times/day, per week, per month? Is change done on a fixed schedule or on an as needed basis? How long does a color change take? Does the time required vary according to the "direction" of the change (S→U, S→TS)? How long does it take to ready the system for the color change -- is there a "slowdown period" involved? How long does it take to achieve "normal" system activity once the new level has been established? How many operators are needed?

6. Briefly describe the steps involved in the color change process.

7. Has the cost of an hour of WWMCCS time ever been established? Typically, commercial rates of $450 to $750 per hour have been estimated for H6000 series machines. Is there a comparable cost for your site?

8. Will future operational requirements necessitate a multilevel processing capability?

Site Survey

Table VII-1 summarizes the WWMCCS sites surveyed and the method of data collection used (e.g., Questionnaire, Telephone or AUTODIN). The order of the sites in the table reflects the chronological sequence of sites surveyed.

The following sections summarize the data collected from each of the surveyed sites.

Table VII-I

Data Collection Site Summary

	WWMCCS Site	In Person Questionnaire	Telephone	AUTODIN	Data Not Avail.
AFDSC	Data Services Center	X			
SAC	Strategic Air Command	X			
MAC	Military Airlift Command	X			
DSDC	Data Systems Design Center		X		
NORAD/ CONAD	Cheyenne Mountain Complex			X	
TAC	Tactical Air Command			X	
ATC	Air Training Command			X	
AU	Air University			X	
PACAF	Pacific Air Command			X	X
USAFE	Air Force Europe			X	X
REDCOM	Ready Command			X	X
AFSC	System Center			X	X

1. AIR FORCE DATA SERVICES CENTER (AFDSC)

On 17 May, 1974, MITRE visited the Air Force Data Services Center (AFDSC) in the Pentagon.

AFDSC WWMCCS SITE OPERATION

AFDSC's approach to security can be categorized as the "high water mark" technique. Their machine room, which houses an H6060 (a typical WWMCCS processor), in addition to other processors, is physically secure. The machine, the terminal areas, the system staff, and the general user population are cleared to the highest security level handled by the system: Top Secret. All programs and their resultant output are processed and handled at this level. Although the mix of jobs is currently 10% Top Secret (TS), 40% Secret (S), and 50% Unclassified (U), the system is always kept in TS mode. Declassification of input and/or output data is left to the discretion of the user. Before downgrading any output, it is the user's responsibility to scan the generated output for inclusion of sensitive information (i.e., dumps, improperly burst printer output, etc.) that has somehow "crept" into his data.

Since this WWMCCS system is dedicated solely to Top Secret processing, no color changes are necessary; and hence, no color changes are performed. AFDSC has chosen the "high water mark" approach to be their site's specific solution to the problem of processing. Classified processing is feasible in their case, due to the availability of a variety of machines that can be used to handle various levels of data. The multi-system environment makes such an approach economically feasible.

Since the primary factor in the operation of a WWMCCS system is responsiveness in a command and control environment, AFDSC considers the color change process to be sufficiently time-consuming to be deemed operationally unacceptable. However, given the existence of automated procedures providing a 5-10 minute color change process, AFDSC would seriously consider revamping its procedures in order to implement the proposed JSS. The new system would, most certainly, have to conform to rigid security standards, perform quick changes, and be responsive to command and control requests.

POTENTIAL PROBLEMS

AFDSC utilizes a variety of computer systems (IBM, Honeywell, etc.). In order to increase operational efficiency, optimize space usage, and facilitate reconfiguration, the peripheral devices accessed by these systems have been physically segregated into "farms" of similar devices. Thus, one can find a printer farm, a reader farm, etc., within the confines of AFDSC. The printer farm concept, although specific to AFDSC, may raise a potential problem in the design of the JSS.

Printers within the printer farm may be allocated to various systems on a dynamic basis. The printers dedicated to the WWMCCS system process only Top Secret data, whereas neighboring printers generate printout of varying sensitivity. The physical adjacency of WWMCCS-printers to non-WWMCCS-printers raises questions directly related to the design of a JSS. The frequency of operator-caused error in detaching printout from the wrong printer may increase with the advent of a JSS. Since the JSS will control device reconfiguration, the operators must be made aware of the multiple roles played by certain devices within their farm. Printers may be interchanged among systems and those printers allocated to a particular system may be required to process output of various security levels at different times. That is, operators can no longer expect to remove level A data from printer A, since level A data will be output to different printers during the day.

Another ramification of printer farm operation is the future impact of a speedy color change via the JSS. Even though the JSS attempts to "take the man out of the loop", thereby decreasing the number of manual operations performed per change, it is very likely that an increase in the number of color changes will tax operators more heavily as they perform repetitive procedures at an increased frequency. The use or misuse of a JSS as a productive tool will, in the end, be decided by the operations staff who must contend with problems indigenous to their site.

Hardware

Figure VII-1 shows the H6060 WWMCCS hardware configuration at AFDSC.

Software

1. The operating system installed on the Air Force Data Services Center WWMCCS H6060 computer is the World Wide (WW) 5.3 Software Release as supplied by JTSA and AFDSDC.

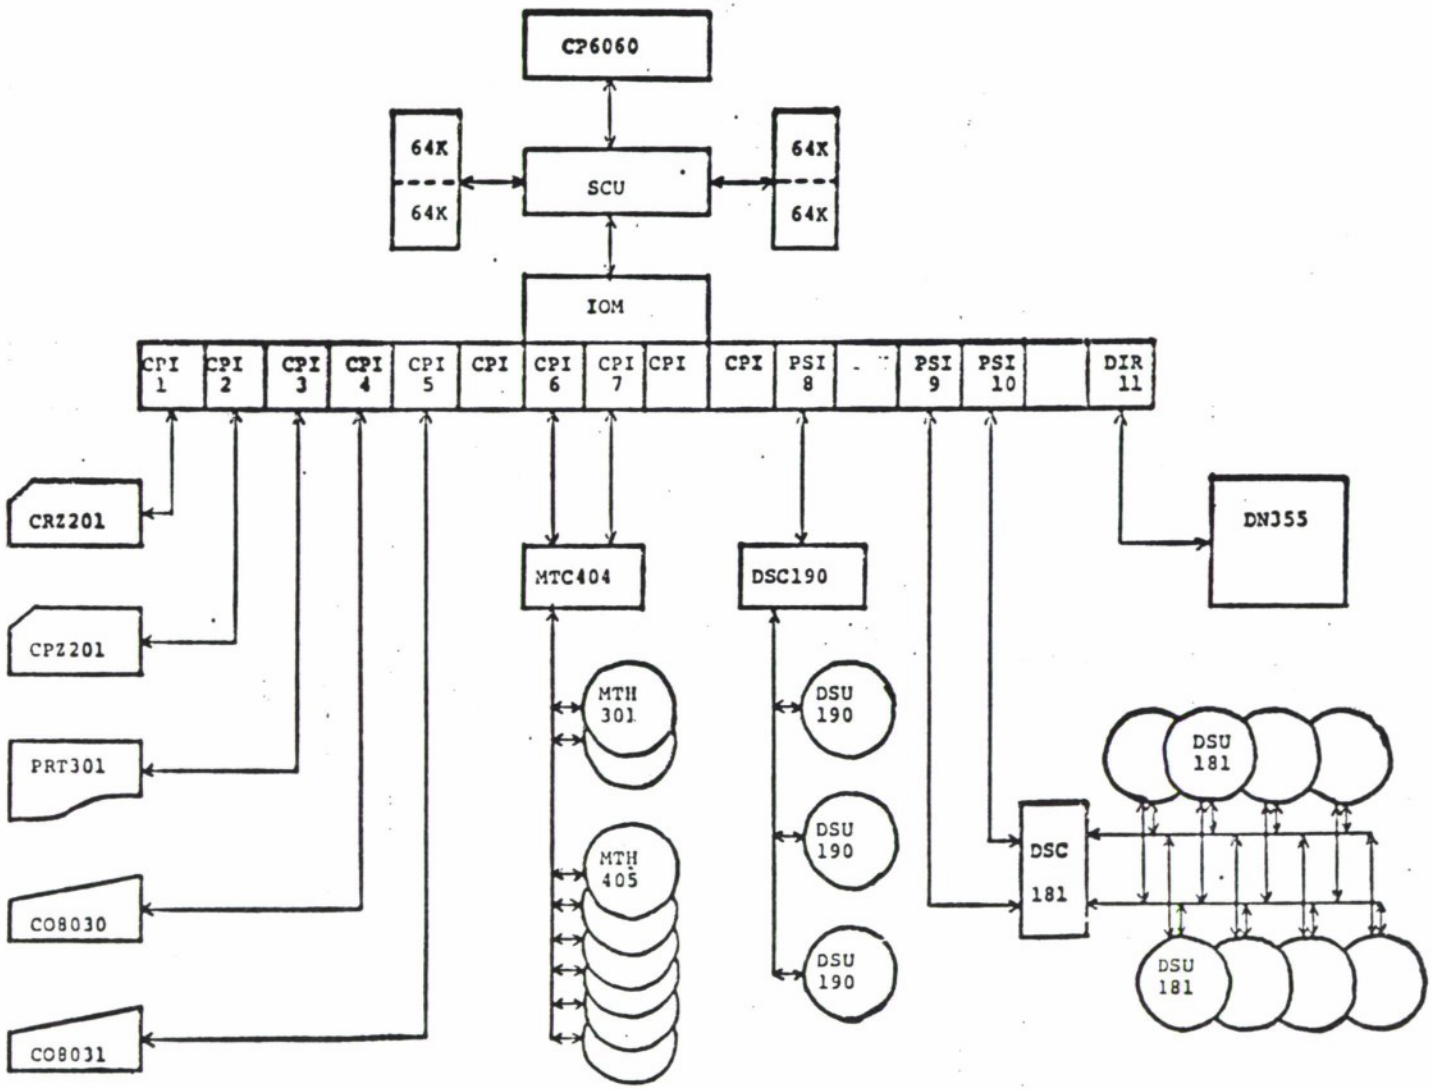

Figure VII-1. AFDSC Planned Configuration

2. Command unique systems software include the following modules:

SCAN - scheduling module
GNAT - Job Identification (ID) card verifier module
XSA2 - table that includes the ID's for GNAT
UMME - output marking module
XSA3 - timer extension module
BRT7 - resources used by activity

3. WWMCCS Standard Applications Software: Airfield Facilities/Great Circle Distance Calculation File.

4. Command Unique Applications Software:

a. USAF Rated Report
b. Key Word Index
c. Status of Forces
d. Frequency Prediction
e. Unit Capability Measuring System
f. Logistics
g. Personnel and Manpower
h. Plans
i. Force Identification
j. Medical
k. Several USAF Standard MAJCOM Systems (Accounting and Finance, Manpower, etc.)

2. STRATEGIC AIR COMMAND (SAC)

On 30 May 1974, MITRE staff visited the Strategic Air Command (SAC) Headquarters at Offutt AFB.

INTRODUCTION TO SAC

In our initial meeting, introductory remarks provided a brief sketch of the history, current operation, and future requirements of SAC's WWMCCS operation from 1971, when the first WWMCCS computer system was installed at SAC. Attention was also given to the netting of their machines with other WWMCCS machines at distant sites.

Brief Concept of Operations:

1. As stated in the SAC WWMCCS Integration Concept the conversion to the WWMCCS standard computer system has occurred over the past several years. The initial equipment installed was composed of two force control systems. This configuration was augmented to provide the capability to process three simultaneous jobstreams in October 1973. This configuration supported the initial conversion phases for NCA reporting, the SAC Automated Command Control System (SACCS) data processing subsystem, Single Integrated Operational Plan (SIOP) production, real-time computational support for the 4000 AEROAG and Major Command (MAJCOM) information processing.

2. A further augmentation was recently implemented to provide the necessary capacity to complete the conversion effort by establishing a fourth jobstream within the WWMCCS configuration. Additional memory, disk storage, and communications are scheduled to support enhancements to the existing systems and normal system growth. Jobs are scheduled across all systems based on current workload and security requirements. The most recently created jobstream supports only batch processing at this time.

SAC's WWMCCS Configuration

The latest augmentation of SAC's configuration established a totally batch jobstream to relieve a saturated workload condition on the existing jobstreams supporting on-line interactive and Remote Job Entry (RJE) requirements. Near term configuration changes will provide an expansion of memory, disk storage and communications capabilities.

SAC currently has two Honeywell dual processor 6080's, one of which has been "split" into two single processors (see Figure VII-2). The three systems are labelled:

1A - uniprocessor 6080
1B - uniprocessor 6080
2 - dual processor 6080

WWMCCS WORKLOAD SHARING

System 1A - On-Line

Force Control On-Line (Unclassified→TS SIOP-ESI).
Force Control Interactive (Unclassified→TS SIOP-ESI).
4000th AEROAG On-Line (Unclassified).

System 1B - Period Processing

Force Control Programming (Unclassified→TS SIOP-ESI).
4000th AEROAG Programming (Unlcassified).
MAJCOM Applications (Unclassified→Top Secret).

System 2 - Period Processing

SIOP Production (Unclassified→TS SIOP-ESI).
SIOP Development (Unclassified→TS SIOP-ESI).
SIOP Programming Support (Unclassified→TS SIOP-ESI).

SAC WWMCCS FUNCTIONS

There are five major functions performed on SAC's WWMCCS machines. None of these are entirely operational on the 6080's - all are in some stage of conversion and/or development. The five functions, according to a brief report given to us by SAC, are:

1. Support of the National Command Authority. Support of the NCA is the primary mission of the SAC WWMCCS system. The command and control and SIOP planning functions at SAC are in direct support of the NCA. As new NCA requirements are identified, they will also be supported on the SAC WWMCCS system.

2. Support of SACCS (SAC's dedicated command and control network). There are three equipment subsystems in the SACCS: Data processing, data transmission, and data display.

	SYSTEM 1A	SYSTEM 1B	SYSTEM 2
Processor	1 CPU	1 CPU	2 CPU
Memory	4 @ 64K each	4 @ 64K each	4 @ 64K each
181 Disk	8	8	10
190 Disk	5	5	10
Tapes	7	7	12
Data Net 355	1	1	2
Printers	1 shared between 1A and 1B		1
	4 Printers available to 1A, 1B, or 2 by Switching		
Readers	1	1	2
Punches	1 shared between 1A and 1B		1
Console	1		2
	1 shared between 1A and 1B		

Figure VII-2. SAC Hardware Configuration

a. Data Processing Subsystem (DPS). The DPS consists of System 1A backed up by Systems 1B and 2.

b. Data Transmission Subsystem (DTS). The DTS provides the communications lines, message switching, and transmission equipment for the entire system. It ties the WWMCCS computers, the Data Display Centrals (DDCs), and the bases having SACCS input and output equipment into one network.

c. Data Display Subsystem (DDS).

3. Support of the 4000th Aerospace Applications Group. The 4000th AEROAG requires real-time, interactive, and off-line remote batch data processing support. The real-time and time-shared processing requirements are supported on the force control on-line system while the batch processing requirement is supported on an off-line system.

4. Development and production of the SIOP. The war gaming and analysis function is vital to the evolution and improvement of the SIOP. However, because of high core and processor time requirements, the extremely large mathematical models necessary for gaming and SIOP are normally processed with a low urgency during periods of relatively low use. Remote entry devices are available during prime time for simulation analysts to run minor excursions and to set up models for later execution.

5. MAJCOM Management Information System Applications. Periods of batch mode data processing on an off-line system are dedicated to supporting the management information requirements of the SAC Headquarters Staff. The primary staff agencies supported are: Logistics, Personnel, Plans, and the Comptroller.

CURRENT SAC WWMCCS USAGE

The following is a breakdown of WWMCCS system usage by processor. The functions described are either currently operational or under development. Note that according to current operational policy, color changes will be done only on systems 1B and 2.

System 1A

Handles on-line message processing, in support of various networking requirements levied on SAC. This set of requirements includes command and control. System 1A will also be used to process unclassified queries from the aerospace group. No color changes will be performed, since this system must be on-line at all times to receive force control traffic and support command posts.

System 1B

Supports the aerospace group and MAJCOM applications. In a crisis situation, it can be used in "shadow mode" to System 1A, duplicating all actions performed by 1A. When System 1A goes down for preventive maintenance (PM), System 1B is configured to run in shadow mode to System 1A, until 1B is capable of carrying 1A's complete workload. The process may be reversed between 1B and 1A. The switchover from 1A to 1B can also be accomplished by stopping 1A, removing 1A's packs and physically remounting them on drives connected to 1B. Alternatively, the crossover can be performed by reconfiguring the IOM's controlling the drives in question, rather than moving packs from drive to drive.

System 2

Can be used as a backup to System 1A. However, its primary function is to support the SIOP development.

SOFTWARE

System software supporting WWMCCS and the local command includes:

1. WWMCCS Standard Systems Software: All WWMCCS Standard system Software released for implementation to date is being utilized at SAC with the exception of the Transaction Processing Executive (TPE).

2. In order to meet the specialized interfaces required to support on-line system requirements, the following system software modules were developed under the 436M Contract;

a. SAC On-Line Interactive Controller (SONIC): Implemented as a sub-executive under GCOS. Interacts with GCOS to provide required on-line priority control.

b. GNTM: DATANET 355 module developed to communicate with the non-WWMCCS CRTs procured under the 436M Contract.

c. GRTM: A module for the DATANET 355 to support communication of several logical lines from the Remote Terminal Facility (RTF) to the DATANET 355 over one physical 38.4K bps encrypted line.

WWMCCS Standard Application Software

All standard application software released for implementation is being utilized at SAC with the exception of the Worldwide Data Management System (WWDMS). The SAC requirement for an on-line DMS is being met through the use of the Force Management Information System (FMIS) and is expected to be operational in the first half of 1975.

Command Unique Applications Software

A great deal of the software converted to the WWMCCS Standard Computer Systems at SAC is unique to command requirements.

OPERATIONAL STATUS OF SAC WWMCCS SITE

The Computer Sciences Corp (CSC), under Air Force Project 436M, is currently directing its efforts to (1) provide the necessary connections between SAC's WWMCCS machines and other SAC systems and networks, and (2) establish a switchable interface between user terminals and each SAC WWMCCS machine. The main product of this effort is a large reconfiguration panel through which all terminals and outside interfaces are connected to one of the three WWMCCS machines. This panel enables a security operations officer to physically connect/disconnect terminals to the WWMCCS complex. This "hands-on" approach allows for the immediate disconnect of malfunctioning or suspect terminals, and constitutes a substantial portion of the SAC security assurance.

As stated above, none of the five main WWMCCS functions (applications) are completely operational at this time. Numerous color changes are currently performed because of ongoing software and hardware development work. The number of color changes will probably be reduced once these functional subsystems become operational. However, because of the volume and variety of processing involved, color changes will still remain a problem for the SAC WWMCCS machines.

COLOR CHANGE SCHEDULING

The current policy is to do color changes by fixed schedule. Each Wednesday, there is a meeting to determine the processing schedule for the upcoming week. In some cases, the final determination as to which group or application is allocated a particular time slot cannot be reached at an operational level and has to be decided at a higher managerial level. This situation is exacerbated by the large amount of development work that is still being done. It is expected that this scheduling procedure will be continued, even after all WWMCCS functions are operational.

COLOR CHANGE PROCEDURE

On-line users are notified of an impending change and are requested to save their files and terminate processing immediately or else suffer abrupt program termination. When a "scheduled" rather than "crisis provoked" color change is imminent, the operators scan the current job queue and submit only those jobs that require minimal execution time and spoolout resources. The decision as to which job within a job class is run is based on data

supplied by the user on the job card. This user-supplied job classification/description policy assumes the trustworthy and knowledgeable user, who can estimate his job's maximum needs. A latent pitfall inherent in this policy becomes evident during the slowdown (job draining) phase, when grossly incorrect user estimates cause unexpected delays in the color change process. It is then up to the operator to compare actual job statistics with user specified maximums, and to terminate the program if estimates have been exceeded.

Occasionally an exigency arises that prevents use of lengthy, time-consuming operations in favor of quick response to higher priority requirements. This situation precludes performing any steps in the color change process geared toward protection of system or user files, or the preservation of system internals that would normally facilitate restoration of the system to its current operational status at a later time. When such a case occurs, jobs in execution may be aborted, and partial spoolout destroyed. The general user's only alternative to loss of data, in the advent of emergency actions undertaken by the system staff, is to submit the job as "re-runable"; the system staff will attempt to rerun those jobs so indicated. The user must, of course, determine if his job can be re-executed.

In either case, whether scheduled or random color changes are pending, the input job queue remains intact at the time the processor is purged and cleared. The job stream will, at a later time, be recreated when the system is processing data at the appropriate security level.

ACTIONS PERFORMED DURING COLOR-CHANGE

The actions taken during a color change depend upon the relative direction of the change in security levels. That is, going to a higher classification (S→TS) involves less chance of security compromise than would going to a lower level (TS→S), and hence, requires fewer actions to be performed.

Four general color changing actions may take place, in varying degrees and order, before a new classification level can be established by booting a new operating system. These actions, described in greater detail in the ensuing discussion are:

1. saving or backing up the current files;
2. physical changing of disk packs and tapes;
3. runout of print queues; and
4. purging or sanitizing the system.

If all four actions are taken, two or three hours can elapse in changing processing levels; a minimum of two hours is normally allotted in the schedule. A change from a lower to a higher level of security, or from one system to another at the same level, may take only 30 minutes if only pack changes are involved.

1. File Saving

SAC personnel consider the file saving procedures as part of the color change process. This inclusion is one of the reasons why the duration of a change is so much greater than normally expected. All files are currently saved by writing them from disk to tape. (Eventually, at least on the on-line system, all saves will be done from disk to disk). At present, it takes from 1 to 1½ hours to save all system files.

2. Changing Packs and Tapes

Physical removal of the various media is required whenever the processor level or type of processing is changed. The operators dismount the tapes and disk packs, put them in carts, wheel these carts out of the machine room, and store the media before preparing to mount the tapes and disk packs for the new level. This is a time-consuming process. As previously noted, however, the physical separation of data by classification and category is an important and carefully enforced concept at SAC.

3. Run-Out of Print Queues

Printer run-out or draining of print queues is also a time-consuming process at SAC, since many of the jobs require large amounts of printing. During a color change, these print jobs can be terminated if necessary, but they are usually allowed to complete.

Systems 1A and 1B can each be configured with up to four printers; System 2 can include up to five. There are currently only seven printers available, however, and one or more of these is often down.

Some printing is occasionally done off-line by transferring print queues to tape and processing these tapes on a stand-alone IBM 1460. However, the printers on the 1460 have a maximum print speed of only 600 LPM and, with the character conversion that must take place, the effective rate is little more than half the rated speed. The SAC operations staff is not enthusiastic about the potential of this alternative because of the limited printer speed of the 1460's and the space problems that would be caused by introducing a number of 1460-like machines to do the off-line printing. SAC has instituted a plan to reduce color change duration by overlapping the run-out and system save processes. However, the operations staff must often wait for print queues to drain after the system is ready to be purged.

4. Purging the System

Purging or sanitizing the system is performed whenever going from a higher to a lower classification. It is usually not required when remaining at the same level or when going to a higher level. Purging involves the scheduling of programs to clear buffers, controllers, memory, etc. In some cases (printers), manual purging is performed by manipulating switches on the control panel. To purge and color change without saving files, once print queues have been emptied, takes approximately 45 minutes.

PROBLEM AREAS

The SAC staff has enumerated several major problem areas, reflecting the current status of their WWMCCS system, that must be taken into account in the design and future implementation of the JSS. These factors are:

Space Limitation

As in most computer centers, physical space limitation imposes a significant burden on the SAC computer system, especially when the need for additional hardware arises. At present, the SAC WWMCCS machine room is teeming with equipment and "moving walls" to provide more space is almost a daily occurrence. Since the JSS will require a minicomputer and possibly a few additional disk drives, the lack of space should be taken into account during the JSS design phase.

Meeting Expanding Processing Requirements

The operations staff also states that, even though they are not yet operational, the maximum processing capacity of the three 6080's is rapidly being neared. This situation, coupled with the inordinate delay associated with color changes, complicates the weekly task of scheduling machine time for various developmental and operational systems currently active. The lack of computing power and processing time, which is a direct result of the waste of resources incurred by changing colors, forces the operations staff to stringently adhere to schedules that are based on requirements and data already a week old.

File Saving Process and Spoolout Queue Run-Out

The saving of files (system and user) prior to a color change is obviously a very time-consuming operation. The save process (disk to tape) entails copying the entire contents of a given disk pack to tape. Rather than limiting the quantity of data saved, emphasis is currently directed toward faster disk-to-tape, and possibly disk-to-disk operations. The implementation of a COM (computer output microfilm) processor is also under consideration as a vehicle for handling vast quantities of printed output. An alternative approach would be a stratified, or selective, file saving technique (similar to the backup utility, BACKUP.SYSDAEMON, used on Multics), whereby only those files modified during the course of operation at the current security level are saved. This technique saves time by writing a smaller number of records to tape, and allows for the creation of a daily master merge tape during processing periods.

ADDITIONAL QUESTIONS

In addition to the problems already described, the following questions arose:

1. Where will the JSS be attached?

Initial proposals concerning the JSS have only superficially broached the issue of actual interconnection of the control mini with the Honeywell processor. The problem faced at SAC is complex because of the multiplicity of systems in operation (i.e., system 1A, 1B, 2). One must consider whether three JSS systems are needed, or only one interchangeable, supervisory JSS.

2. What effect would a 5-10 minute color-change have on operators? On scheduling?

The future role of SAC operators as users of the JSS remains undefined. The actual benefits of the JSS may be obscured, at first, until stable operational procedures have been established via hands-on experience.

3. What is the minimum configuration that can define the start of a new processing level (i.e., how many disks, tape, etc.)?

The concept of dynamic reconfiguration of peripherals may be introduced as part of the JSS. This concept implies that a minimum number of peripherals will be brought up initially, when a new level begins processing. Additional units will be made available as soon as the operator dismounts the old media, remounts the new, and notifies the system of his action. Thus, there will be some delay before all available peripherals can be used at a new level.

It was difficult to determine the minimum number of disks, tapes, etc., that are necessary to start effective processing when a new system is brought up. At SAC, parts of GCOS are spread across a number of the system disks in an attempt to improve system performance. This technique might be modified with the advent of the JSS, so that the GCOS operating system would be fit onto a smaller number of disks. Degradation in system performance could result from this redistribution of GCOS.

3. MILITARY AIRLIFT COMMAND (MAC)

On June 20, 1974, MITRE staff visited Military Airlift Command (MAC) Headquarters at Scott AFB, Illinois.

ABSTRACT

MAC's interest in reducing computer dead-time associated with the color changing process coincides favorably with the ESD/MITRE effort to design a Jobstream Separator.

The recent Israeli conflict spurred MAC to ferret out weakpoints in its operation during crisis mode. MAC has proposed procurement of an additional 6080 processor that would improve ADP operations when processing classified data. This dedicated processor would reduce the possibility of security violation in processing multi-level information.

MAC's staff expressed a definite interest in the JSS, especially in light of operations to be performed on their new system. The staff willingly took several hours to answer, in detail, the questionnaire. A color change from unclassified to Secret level was demonstrated. Lastly, several interesting observations and valuable suggestions were made by the operations staff that should aid in the design of the JSS.

MAC OVERVIEW

The Military Airlift Command (MAC) mission is to provide airlift services for the transportation of passengers and cargo - the global deployment of units, and special airlifts for purposes of resupply during peace and contingency operations. MAC must support daily DoD and Service requirements and, in addition, must respond to worldwide contingencies on a 24 hour-a-day, seven day-a-week basis.

MAC OPERATIONS

MAC has special operational problems (distinct from automation problems) that arise during crisis situations. MAC's data base contains predominantly unclassified information that can be accessed from dozens of non-secure terminals. During emergency operation specific portions of the data base that pertain to the crisis-provoking situation become classified. The mixture of security levels results in the contamination of the entire data base. The "tainted" (Secret or Top-Secret) portions of the data must, therefore, be accessed by a secure system during a contingency. This requirement results in the disruption of normal (Unclassified) processing and the termination of on-line interactive processing.

The introduction of Secret and Top Secret data into the Unclassified data base occurs under the following conditions:

1. Top Secret Data is introduced in support of REDCOM requirements during contingency conditions and war game planning.

2. Secret Data is introduced when certain application programs are run. For example, the FORSTAT system (Force Control) is run every evening at the same time. Immediate response time is not the key issue in the creation of a Secret processing environment for FORSTAT. However, the time lost in changing colors prevents optimum system utilization, especially since interactive users are denied access to the system when it is in SECRET mode.

Other processing problems such as expanding processing demands generated by an ever growing user population, coupled with limited processing resources, make scheduling the usage of the WWMCCS machines extremely difficult. Allocation of system resources to various offices has created difficulties for MAC's operations staff.

MAC's operation also poses a problem that is contrary to SAC's operation; MAC has a large population of non-secure terminals, thus necessitating the exclusion of terminal users during classified processing. Although current figures show a relative increase in the volume of classified processing to a new high (approximately 25% of normal workload), these terminals play a significant part in shaping MAC's operational and security related policies. MAC's mission must ultimately support in excess of 100 non-secure terminals.

OPERATIONS

Our conversation with MAC operations included the following topics:

1. The requirements levied on MAC by REDCOM and JCS.

2. The operational policies designed to meet these requirements.

3. Operational problems that arise during emergency situations.

4. Proposed solutions to these problems.

Crisis Management

MAC is responsible to REDCOM, which is, in turn, responsible to the JCS. The current MAC configuration has a split 6080, one half of which is operated in classified mode on a 4-6 hour (per level) schedule. Advanced planning programs are run to keep logistical tags on people and equipment to be moved. When a crisis occurs, the JCS needs to know how fast material can be moved and, in addition, needs alternative flow analysis reports concerning aircraft/airlift support. MAC must determine if it can perform the task outlined by REDCOM. A 30-minute limitation placed on REDCOM by the JCS imposes a severe restriction on MAC operations, since it inherits this limit via REDCOM requirements. It takes approximately 30 minutes to perform a color change from Unclassified to Secret levels, plus several hours, typically, to run the analysis programs at the new level; MAC therefore cannot fulfill its processing obligations within the time-span specified by the JCS. If the time needed by MAC to change colors were reduced to 5 minutes via the introduction of a JSS, then MAC could at least begin processing, in order to perform its assigned mission.

Additional H6080 Processor

MAC has proposed, after extensive analysis of its operations, to procure an additional H6080 processor in order to meet REDCOM requirements. Figure VII-3 indicates the current MAC configuration with future alterations indicated by underlines. The use of the symbols "1", "2", and "3" to indicate processor is intended to clarify the drawing - they are not MAC designators.

The existing dual 6080 has been configured (split) into 2 single 6080s, hereby designated System 1 and System 2. System 1 will acquire 128K more core, and System 2, 192K more core; the total core for each system, 256K. The proposed new processor, System 3, will have 192K core, its own DN355, multiplexor and disks.

Systems 1 and 2, which comprise the unclassified force control system, will ultimately be recombined into a single processor when System 3 becomes operational. Since System 3 will be required to handle classified processing, which represents a small portion of the normal workload, it will also be used to lessen the Unclassified workload by handling the MAJCOM update application in a tentative classified mode. Approximately 30 terminals, in 8 or 9 locations, will be attached to System 3 via encryption devices.

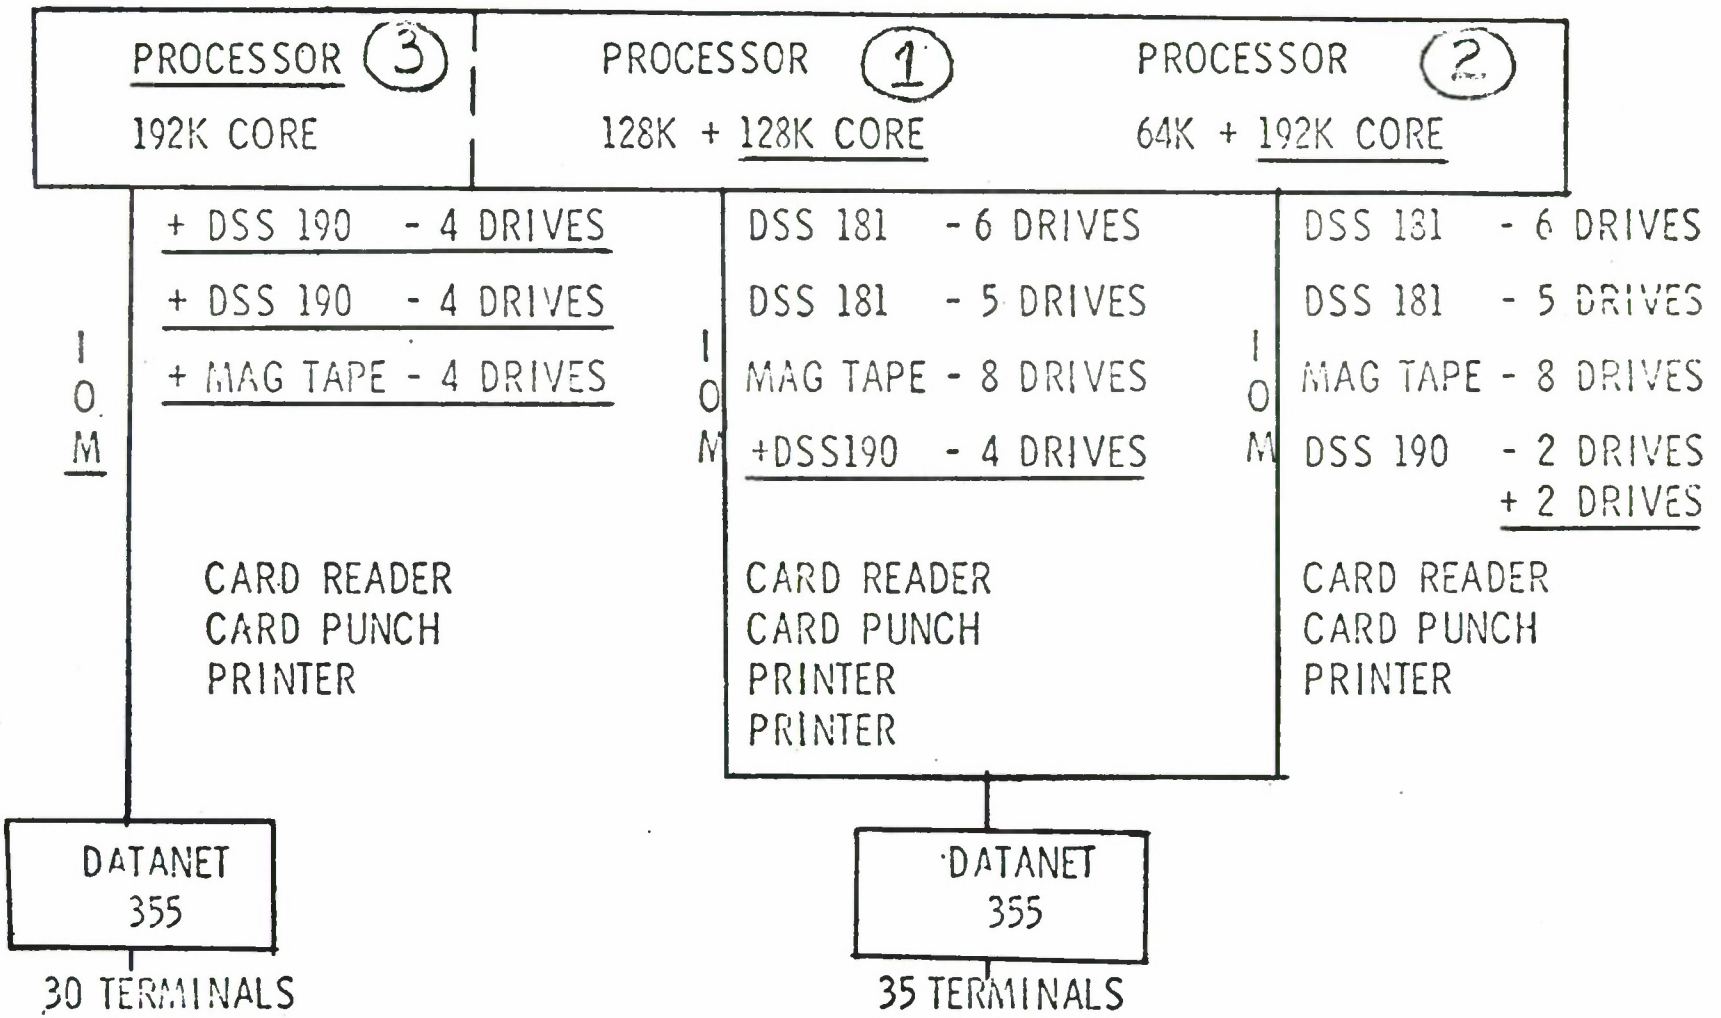

Figure VII-3. HQ MAC Honeywell 6080 Computer System Enhanced Split Configuration

System 3 Operation

Should a crisis arise, any jobs currently in execution on System 3 will be summarily aborted. The Force Control data base residing on disk packs used by Systems 1 and 2 will be copied (disk-to-disk) and physically transferred to System 3, where it can be used for contingency planning. Since System 3 is mainly responsible for handling data concerning resources involved in a contingency, and has no means of accepting unclassified data describing resources not involved in the crisis, periodic updates to the classified version of the data base (System 3) must be made. Initially, System 3 has an up-to-date version of the unclassified data base, but as time progresses, portions of the data base become out-of-date. Unclassified on-line users on Systems 1 and 2 may be interrupted (logged-off) during the disk copy process, so that users cannot tamper with data immediately prior to or during the (update) preparation process.

Color Change Process

There are usually 2 changes made per day: U⟶S and S⟶U. Three times per week the schedule becomes U⟶S⟶TS⟶U. The following is a breakdown of the color change process by phase:

a. Pre-change - (Draining period) - Starts 1 hour prior to scheduled color change. Jobs in job queue are selected by the operator in order to maximize system throughput -- decrease the number of jobs awaiting execution by selecting those jobs that require minimal CPU time and system resources. SYSOUT is copied to tape for later printing. Terminal users are locked out one half hour before change (NCALL function issued by operator). Users are notified of upcoming color change by consulting preset schedule. A courtesy phone call by the operations staff officer, however, usually precedes terminal disconnect.

b. During Change - Requires approximately 15 minutes to perform steps enumerated in checklist. Switch is thrown on device near console to disable phone lines at junction box connection rather than at modem. All devices are cleared (SC units, DN-355's, IOMs device controllers, memory, etc.). At least 90% of this period consists of manual operation.

c. Post-change - Lasts approximately 15 minutes. Old system has been purged, new system is loaded, devices brought up to speed.

Average Workload Statistics

1. Interactive jobs per day: 200 unclassified

   Batch: 200 Unclassified
          20 Secret
           5 Top Secret

2. Mounts/dismounts per level

          U 100 tape mounts, 50 disk
          S 40 tape mounts
         TS 40 tape mounts

3. Number of packs per classification

          U 29 packs, 4 contain GCOS
          S 11 packs, 4 contain GCOS
         TS   8 packs, 4 contain GCOS

EYEWITNESS ACCOUNT OF COLOR CHANGE PROCEDURE

This section presents observations on the salient aspects of MAC's color change procedure. Suggestions and comments concerning current and future operational procedures are included in this section.

Color Change Checklist

Prior to changing colors, a checklist detailing the operations to be performed is produced on the high-speed printer.

Once the system has been quiesced (i.e., there are no jobs left to run in the system), a color change is performed. Typical manual operations performed during the change are: setting and resetting switches on the inside panels of various devices (i.e., 2 Datanet 355's, 4 system control units, device controllers, multiplexors, etc.), dismounting disk packs of old color, color coding the new packs, changing printer and console ribbons, changing card decks, and even changing jackets worn by operators to signify the sensitivity of the new system coming up. Modems are cleared by the process of "bussing-back", whereby the input and output lines are connected. Of the total half hour duration of a color change, approximately 60% is spent performing manual operations under a check-and-double-check surveillance system. This approximation includes the time required for a disk to be powered down, unloaded and remounted. It does not account for the period during which the disk heads are brush-cleaned and allowed to "come-up-to-speed" (approximately 1 1/2 minutes per drive.

Home-Made Phone Disconnect Device

The most unusual hardware feature noted during the tour of MAC's WWMCCS computer center was a "home-made contraption" mounted on the side of the operator's console. The small device, containing approximately 20 pencil-thin switches, was built in-house for a cost of $50.00. The device is used by the operator in order to cut off all terminal users when classified processing is initiated. The device grounds out the phone lines at the junction box, disabling them in a matter of seconds. Prior to the existence of this centrally located switching facility, operators were required to turn off connections at modems located on shelves some 50 feet from the operator's console. The device is indicative of the willingness of the ADP staff to use whatever options are available to them in order to save time and/or effort.

Automating Manual Operations Via JSS

By far, the most time-consuming operation of the manual procedures employed is the setting, resetting, and double-checking of the dense population of switches located on the inside panels of the various controllers. Significant emphasis should be placed on determining the feasibility of clearing the numerous switches in question via a JSS or a JSS-driven reconfiguration panel. The importance of such a facility is quite clear - it is estimated that 10-15 minutes can be saved in performing manual operations.

4. DATA SYSTEMS DESIGN CENTER (DSDC) AND AIR UNIVERSITY (AU)

On August 16, 1974, a telephone call was made to the Data Systems Design Center (DSDC) at Gunter AFB in Montgomery, Alabama, for the purpose of collecting data for the Jobstream Separator analysis. Since the JSS may eventually serve the entire AFWWMCCS community, it is important that MITRE's feasibility study sample the WWMCCS sites that employ smaller, single-processor H6000 configurations to handle the relatively light mix of classified and unclassified data representative of their workloads.

DSDC Mission

The H6000 computer is located in the AF Data System Design Center, Gunter AFB, AL, with a H720 located in the Air University Data Automation Center, Bldg 910, Maxwell AFB, AL. The mission of the Air University is military education, while the mission of the AFDSDC is designing and programming of standard Air Force systems. The H6000 computer is used for AFALT Support, for AU Unique Systems, for AFDSDC Test and Development activities, and the Air University Educational Time-Sharing System. The Air University Unique Systems, which are designed and programmed by Data Automation, support programs of the resident and non-resident schools, i.e., Air War College, Air Command and Staff College, Squadron Officers School, Extension Course Institute, AFROTC, Air Force Institute of Technology, etc., of the Air University. The AU Educational Time-Sharing System provides support to the professional military education classroom exercises, whereby all students have the opportunity to participate and gain experience in the use of computers. The remotes at the various schools operate on-line approximately 9 hours per day, 5 days a week. The other remote terminals and the H720 computers at Maxwell AFB and Gunter AFB are operated on-line. At present, AU does not have the software nor the hardware to operate off-line. Cards are transmitted to the H6000 in the batch mode.

DSDC ADP Equipment

1	CS6060S	Central System incl. processor, system controller, 96K Memory
1	MM6050	Extra Memory, 32K
2	MM6061W	Extra Memory, 64K
2	DSS 181	Disk Storage System (incl 3 DSU 181 each)
10	DSU 181	Disk Storage Units
2	DCH 181	Dual Sim. Channels
1	MTC 404	Magnetic Tape Controller
2	MTH 301	Magnetic Tape Handlers 7 track
11	MTH 405	Magnetic Tape Handlers 8 track
1	CRZ 201	Card Reader
1	CPZ 201	Card Punch
2	PRT 201	Printers
1	SPB 355	Datanet 355
1	RLP 300	Remote Line Printer
1	VIP 786W	CRT
1	KSR 33B	TTY (188C)

Communications

1	9600 Baud Circuit between GAFS, and Maxwell AFB (MAFB)
1	Cross-town Circuit between GAFS, and MAFB
30	Cross-town Circuits between GAFS and Air University school locations, MAFB

Air Force owned and operated base communication lines are used for all other connections between all other on-base devices.

Operations

The data processed at DSDC is predominantly unclassified -- approximately 99% of the data bases is used to develop and test standard system software and to solve the Air University's classroom assignments. The single classified (Secret) task handled by DSDC that prompts the performance of a color change, is the Air Force Alternate (AFALT) program -- a program that re-creates an alternate (back-up) Pentagon command center in the advent of a direct nuclear attack on the Pentagon. AFALT, which is run once a month at a pre-scheduled time, constitutes the bulk of DSDC's sensitive processing. Future extension of DSDC's remote processing capability, via the utilization of additional satellite processors, may increase classified processing requirements as distant commands tap DSDC resources.

PHYSICAL ENVIRONMENT

DSDC maintains all of its equipment in a single room. All peripheral devices of the same class (tapes, disks, except terminals) are "pooled" together and located in the same room as the processor.

SOFTWARE

No modifications have been made to either GCOS or GRTS. DSDC has researched the possibility of using a previous release of Honeywell's checkpoint/restart system but has abandoned efforts due to the poor quality and ineffectiveness of the system. Efforts are currently underway to analyze the most recent version of the checkpoint/restart facility released under the latest version of GCOS.

COLOR CHANGE PROCESS

The color-change, which lasts on the average approximately 1/2 hour in either direction, consists of the following standard actions in changing from Unclassified to Secret.

1. Tell users of upcoming change.
2. Remove all users from system -- jobs are usually allowed to complete before color change is initiated due to lack of urgency in performing monthly color change.
3. Shut down H6000 system.
4. Clear core.
5. Mount Special packs.
6. Cordon off processing area and hang "Secret" signs all over.
7. Shut down the DN 355.
8. Disconnect all phone lines.
9. Reload new system.
10. Run at Secret level.

DSDC FUTURE OPERATION

DSDC will soon be required to run the MPS system (a logistical manpower-personnel system), which by the very nature of the data being processed will require periodic processing at the Secret level. It is very likely that a proposed additional processor will be used as a dedicated (Secret) processor.

5. NORTH AMERICAN AIR DEFENSE COMMAND (NORAD) AND CONTINENTAL AIR DEFENSE COMMAND (CONAD)

WWMCCS Facility Function

As part of WWMCCS, the Honeywell 6080 computers support the National Command Authorities (NCA) and the JCS CINCNORAD. The functions are supported by three major systems: the NORAD Computer System (NCS), the Space Computational Center (SCC) and the Ballistic Missile Defense Center (BMDC).

1. The NCS will provide data processing and display capabilities in support of the NORAD Combat Operations Center (NCOC). The major functional areas are: System Control and Support, Warning and Attack Assessment, Defense and Force Control and Reconstitution Processing.

2. The SCC provides the means for CINCNORAD to exercise operational control/command of the assigned space surveillance and tracking systems and space defense weapons systems.

3. The BMDC is the operating center and the primary means for CINCONAD to exercise operational command/control over the SAFEGUARD Ballistic Missile Defense System. The BMDC uses SAFEGUARD designed computer equipment.

ADP Equipment

Currently, there are three Philco 212 computers supported by an Input/Output Data Controller and display subsystem. In the near future, the three systems will be replaced by three Honeywell 6000 series system -- 2 H6080's and 1 H6050, each with two CPUs, three Input/Output Multiplexors (IOM), two System Controllers and 256K core. There will be a peripheral pool of card punches, readers, magnetic tapes and disks for assignment to any of the three nodes. DN 355s will be available for interfacing with VIPs, remote terminals and for internetting with other WWMCCS ADP installations. Dual Data General Nova 800 computers will interface the 6080s with the Communications System Segment (CSS).

NORAD Interfaces

The NORAD Cheyenne Mountain Complex (NCMC) Improvement Program (427M) will acquire new data processing equipment, software, communications, and displays in a carefully engineered system configuration. Under Program 427M, the NORAD Computer System (NCS), the Space Computational Center (SCC) and the Communications System Segment (CSS) will provide the NCMC with an integrated, responsive capability and a growth potential that will meet a projected life span of 10 years without replacement of major equipment or major software changes.

The 427M System will consist of three system segments and two remote facilities. These segments and facilities, composed of government furnished and specially procured elements, will replace certain systems or equipments in the NCMC. They will be integrated with each other and with certain existing systems or new systems to be installed in the NCMC. The three system segments are:

1. The NCS (H6080) which will replace the NORAD Combat Operations System (Philco 212).

2. The SCC (H6080) which will replace the Space Defense Center (SDC) (Philco 212).

3. The NCMC CSS (H6050) which will replace the Automatic Digital Relay (ADR), the Input/Output Data Controller (IODC), and the Channel Technical Control Facility (CTCF).

NORAD's two WWMCCS H 6080 computers will be configured as three nodes. One node will operate as a realtime system using the NORAD Computer System (NCS) software. The second node also operates as a realtime system using the Space Computational Center (SCC) software. The third node will be used to meet overall operational availability and workload requirements.

NORAD Operation

The NORAD Computer System (NCS) provides the real-time command and control system (directed by CINCNORAD) to direct the aerospace defense of North America. The system is currently under development. In the future, application programs will be used in the receipt, processing, display and output of real-time aerospace defense information (e.g., missile warning and attack assessment). Other application programs will process less time-sensitive information in background mode.

Before the NCS formally joins the WWMCCS Intercomputer Network (WIN), NCS will process at Unclassified and Secret levels. After joining WIN, the NCS will operate in a dedicated mode and in a controlled environment as defined by 5200.28. Color changes will not be required. NCS development activity does, however, necessitate a color change one or two times per week. Going Secret is now considered a timely overhead since processing normally begins at the unclassified level.

Future Operation

Future operations at NCS will be at "system high", where all users have access to the entire system since the system contains no security caveats. According to NCS staff, "if NCS had available, a certifiable multilevel security capability, the possibility of shared usage of computer power when real-time activity is low, would be very attractive. If and when we join WIN, a multilevel operation would be highly desirable. It may be that the lack of such a capability will be a major factor in the delay of WIN."

CONTINENTAL AIR DEFENSE (CONAD)

The Continental Air Defense system (CONAD), which is part of the NORAD/CONAD/ADC complex at Cheyenne Mountain, is described below. The data was collected via AUTODIN message.

ADP Equipment

2 standard H 6060 systems each with 256K core
2 IOM (1 for each system)
2 SC (1 for each system)
16 181 Disk Packs (8 for each system)
6 190 Disk Packs (3 for each system)
4 MPC (1 for 181 and 190 for each system)
4 7-track tapes (2 for each system)
6 9-track tapes (3 for each system)
2 DN 355
2 HSLA
2 201 printers (1 for each system)
3 Card Readers (2 system A, 1 System B)
2 Card Punches

Operation

CONAD is fully operational. Standard IDS (Integrated Data Store) and classified CONAD unique programs are run. Multilevel and compartmented information is processed, but only at "system high". Color changes are not performed.

Future

As to the future need for multilevel processing capability, CONAD responded: "We are assuming the multilevel security mode is as defined in paragraph 10, enclosure 2, DoD Directive 5200.28. Our current and future ADP capability would be enhanced if we could operate in a multilevel security mode as defined in 5200.28. We operate at system high because the multilevel security problem has not been solved. We will continue to operate at system high until the multilevel security problem is solved, if ever."

Cost Per Hour

An hour of WWMCCS time on the CONAD processor is estimated to cost approximately $429.00.

6. TACTICAL AIR COMMAND (TAC)

The following is a summary of data collected via AUTODIN from the Tactical Air Command at Langley AFB, Virginia.

Mission/Function

As part of WWMCCS, the TAC H6000 computer supports requirements of the National Command Authorities and the JCS and is responsible to other requirements of the Tactical Air Command. The HQ TAC Command and Control System provides Commander TAC/CINCAFRED an automated command and control support system with an effective capability to plan, manage and control TAC resources in the execution of its prescribed functions during normal and emergency situations and in a compressed time/space environment. The mission of this system is to provide information required to: (1) determine the ability of TAC to execute operational plans; (2) measure the effects of current operations on TAC's ability to support general, small, or counter-insurgency warfare and other supporting operational plans; (3) determine the support required for TAC forces engaged in general or small war, show of force, counter-insurgency, exercises, maneuvers, and rescue operations; (4) make recommendations to USCINCRED or Chief of Staff USAF, concerning optimun utilization of TAC forces and deployment or redeployment of TAC/AFRED forces; (5) furnish USREDCOM and USAF with current status and capability of AFRED forces and other data as specified by these headquarters; (6) assist USREDCOM in the generation of operational plans for contingency, counter-insurgency, and exercise operations; (7) direct the implementation of contingency and emergency operations; (8) monitor force status, commitments, locations, short alert requirements, programming and flying time for more effective utilization and efficient operation.

System Usage

Remote terminal sites to support staff users will be used throughout the life of the system. On-line support through time-sharing will be provided to support certain functional activities. Certain systems will be remote/local batch oriented with emphasis on remote operations. Other systems will operate in a near real-time mode through remote communications and CRT terminals.

TAC WWMCCS ADPE

1	CS6060S	Central System 6060 Central Processor System Controller MOS Memory (128K) Core Memory (256K) I/O Controller Type 1 Master Console EIS Feature for CS6060S
1	DSS181	Disk Storage Subsystem MPC181-1 DSU181-3
5	DSU181	Disk Storage Unit
1	DSS190	Disk Storage Subsystem MPC190-1 DSU190-2
7	DSU190	Disk Storage Unit
1	MTH400	Mag Tape Controller
2	MTC301	Mag Tape 7 Channel
6	MTC405	Mag Tape 9 Channel
2	PRT201	Printer
1	CRZ201	Card Reader
1	CPZ201	Card Punch
3	RLP300	Remote Printer
14	VIP786W	CRT Subsystem
1	KSR33A	Teletype EIA
1	KSR33B	Teletype 188C
1	SPB355	Datanet-355 Processor

TAC Software

1. Force Status and Identity Report (FORSTAT).

Purpose/Function: The FORSTAT system consists of 13 TAC-unique programs with 19 supporting files and tables. The function of the FORSTAT Reporting System is to provide current operational information through the reporting of the status of both personnel and equipment. TAC units report daily all changes in status via AUTODIN.

2. TAC Automated Planning System (TAPS).

Purpose/Function: The TAPS System presently consists of 39 TAC-unique programs which interface with seven files. The system assists the Deputy Chief of Staff for Plans, Operations, Personnel, Civil Engineering, Logistics, Communications/Electronics and the Command Surgeon in the development of deployment planning data in support of Operations Plans as well as a means of rapidly responding to planning requirements levied by HQ USAF, REDCOM and LANTCOM. This system allows the TAC staff to generate a complete OPlan force list and tailor it to fit any contingency.

3. TAC Aircraft Profiler (TACAP).

Purpose/Function: TACAP consists of 15 TAC-unique programs and four supporting files. TACAP is designed to assist the TAC operations planners in the generation of deployment aircraft profiles for rapid response to worldwide commitments.

4. TAC Airlift Management System (TAMS).

Purpose/Function: The TAMS System provides the TAC airlift staff with automated support and consists of three major subsystems: Scheduling, Planning, and Flight Following.

Color Change - Description

Color change procedures are performed at TAC. Unclassified and Secret data/programs are processed in the same job mix (batch) and are considered to be one color. TS is processed approximately three times a week. Color changes are performed from U or S to TS and from TS to U or S. Color changes are done on an as-needed basis and the time required to perform a change varies. If there are no remote terminals on line, it takes two minutes to secure the computer area and to physically disconnect the KG-34's from the remote

devices. If there are remote terminals on line, TS usage must be weighed against the urgency of the pending TS job. In going from TS to S or U, two minutes are required to purge the system, turn on the KG-34's (cryptographic device) and remove barriers from the computer area. No additional operators are needed to perform the color change. According to ADP staff, there is no noticeable processing slowdown associated with color change process.

Color change activity includes the following steps:

Notify all TSS and $TRAX users of intentions of process TS. Issue an "NCALL". Once users are signed off all KG-34's are physically disconnected. Signs are posted on computer room doors notifying personnel that TS is being processed.

7. AIR TRAINING COMMAND (ATC)

The following is a summary of the data collected via TWX from the Air Training Command at Randolph AFB.

Mission -- to provide for WWMCCS training.

Operation -- no classified data is handled, hence no color change is performed.

Cost -- the cost of one hour of CPU time for the H6060 at ATC has been estimated at $398, as per AFM-300-12.

Configuration

1	CS6060S		Central System
1	CP6060		Central Processor
1	SC6060		System Controller
4	MM6060	32K	Memory
1	MM6061	64K	Memory
1	IOC		Input/Output Controller
1	CRZ201		Card reader
1	CPZ20		Card punch
1	PRT201		Printer
1	SPB355		Data Net 355 processor
1	MTH400		Mag tape controller
1	MTC301		Mag tape 7-channel
2	MTC405		Mag tape 9-channel
1	MPC 181		Disks
4	DSU181		Disk storage subsystem
1	Master Console		
6	CRTS		